



Multi-party Threshold Private Set Intersection with Sublinear Communication

Saikrishna Badrinarayanan^{1(✉)}, Peihan Miao², Srinivasan Raghuraman¹,
and Peter Rindal¹

¹ Visa Research, Palo Alto, USA

{sabadrin,srraghur,perindal}@visa.com

² University of Illinois at Chicago, Chicago, USA
peihan@uic.edu

Abstract. In multi-party threshold private set intersection (PSI), n parties each with a private set wish to compute the intersection of their sets if the intersection is sufficiently large. Previously, Ghosh and Simkin (CRYPTO 2019) studied this problem for the two-party case and demonstrated interesting lower and upper bounds on the communication complexity. In this work, we investigate the communication complexity of the multi-party setting ($n \geq 2$). We consider two functionalities for multi-party threshold PSI. In the first, parties learn the intersection if each of their sets and the intersection differ by at most T . In the second functionality, parties learn the intersection if the union of all their sets and the intersection differ by at most T .

For both functionalities, we show that any protocol must have communication complexity $\Omega(nT)$. We build protocols with a matching upper bound of $O(nT)$ communication complexity for both functionalities assuming threshold FHE. We also construct a computationally more efficient protocol for the second functionality with communication complexity $\tilde{O}(nT)$ under a weaker assumption of threshold additive homomorphic encryption. As a direct implication, we solve one of the open problems in the work of Ghosh and Simkin (CRYPTO 2019) by designing a two-party protocol with communication cost $\tilde{O}(T)$ from assumptions weaker than FHE.

As a consequence of our results, we achieve the first “regular” multi-party PSI protocol where the communication complexity only grows with the size of the set difference and does not depend on the size of the input sets.

1 Introduction

Private set intersection (PSI) protocols allow several mutually distrustful parties P_1, P_2, \dots, P_n each holding a private set S_1, S_2, \dots, S_n respectively to jointly compute the intersection $I = \bigcap_{i=1}^n S_i$ without revealing any other information. PSI has numerous privacy-preserving applications, e.g., DNA testing and pattern matching [TPKC07], remote diagnostics [BPSW07], botnet

detection [NMH+10], online advertising [IKN+17, MPR+20]. Over the last years enormous progress has been made towards realizing this functionality efficiently [HFH99, FNP04, KS05, DCT10, DCW13, PSZ14, PSSZ15, KKRT16, OOS16, RR17, KMP+17, HV17, PSWW18, PRTY19, GN19, PRTY20, CM20] in the two-party, multi-party, and server-aided settings with both semi-honest and malicious security.

Threshold PSI. In certain scenarios, the standard PSI functionality is not sufficient. In particular, the parties may only be willing to reveal the intersection if they have a *large* intersection. For example, in privacy-preserving data mining and machine learning [MZ17] where the data is vertically partitioned among multiple parties (that is, each party holds different features of the same object), the parties may want to learn the intersection of their datasets and start their collaboration only if their common dataset is sufficiently large. If their common dataset is too small, in which case they are not interested in collaboration, it is undesirable to let them learn the intersection. In privacy-preserving ride sharing [HOS17], multiple users only want to share a ride if large parts of their trajectories on a map intersect. In this case, the users may be interested in the intersection of their routes, but only when the intersection is large. This problem can be formalized as *threshold private set intersection*, where, roughly speaking, the parties only learn the intersection if their sets differ by at most T elements.

Many works [FNP04, HOS17, PSWW18, ZC18, GN19] achieve this functionality by first computing the cardinality of the intersection and then checking if this is sufficiently large. The communication complexity of these protocols scales at least linearly in the size of the smallest input set. Notice that Freedman et al. [FNP04] proved a lower bound of $\Omega(m)$ on the communication complexity of any private set intersection protocol, where m is the size of the smallest input set. This lower bound directly extends to protocols that only compute the cardinality of the intersection, which constitutes a fundamental barrier to the efficiency of the above protocols.

Recently, the beautiful work of Ghosh and Simkin [GS19a] revisited the communication complexity of *two-party* threshold PSI and demonstrated that the $\Omega(m)$ lower bound can be circumvented by performing a private intersection cardinality testing (i.e., testing whether the intersection is sufficiently large) instead of computing the actual cardinality. After passing the cardinality testing, their protocol allows each party to learn the set difference, where the communication complexity only grows with T , which could be sublinear in m . Specifically, [GS19a] proved a communication lower bound of $\Omega(T)$ for two-party threshold PSI and presented a protocol achieving a matching upper bound $O(T)$ based on fully homomorphic encryption (FHE). They also showed a computationally more efficient protocol with communication complexity of $\widetilde{O}(T^2)$ based on weaker assumptions, namely additively homomorphic encryption (AHE).

In this work, we investigate the communication complexity of *multi-party* threshold PSI. In particular, we ask the question of whether sublinear lower and upper bounds can also be achieved in the multi-party setting.

1.1 Our Contributions

We first identify and formalize the definition of multi-party threshold private set intersection. We put forth and study *two* functionalities that are in fact equivalent in the two-party case but are vastly different in the multi-party scenario. Assume there are n parties P_1, P_2, \dots, P_n , and each party P_i holds a private set S_i of size m . The first functionality allows the parties to learn the intersection $I = \bigcap_{i=1}^n S_i$ only if $\forall i, |S_i \setminus I| \leq T$, or equivalently, $|I| \geq m - T$. In the second functionality, the parties can learn the intersection I only if $|(\bigcup_{i=1}^n S_i) \setminus I| \leq T$.

We briefly discuss the difference between the two functionalities. The first functionality focuses on whether the intersection is sufficiently large, hence we call it $\mathcal{F}_{\text{TPSI-int}}$. The second functionality focuses on whether the set difference is sufficiently small, thus we call it $\mathcal{F}_{\text{TPSI-diff}}$. In the two-party case, we have the guarantee that $|(\bigcup_{i=1}^n S_i) \setminus I| = 2 \cdot |S_i \setminus I|$, so we do *not* have to differentiate between these two functionalities. However, in the multi-party case, we only know that $2 \cdot |S_i \setminus I| \leq |(\bigcup_{i=1}^n S_i) \setminus I| \leq n \cdot |S_i \setminus I|$, hence the two functionalities could lead to very different outcomes. Which functionality to choose and what threshold to set in practice highly depend on the actual application.

Sublinear Communication. The core contribution of this work is demonstrating sublinear (in the set sizes) communication lower and upper bounds for both functionalities. We summarize our results in Table 1. For lower bound, we prove that both functionalities require at least $\Omega(nT)$ bits of communication. For upper bound, we present protocols for both functionalities achieving a matching upper bound of $O(nT)$ based on n -out-of- n threshold fully homomorphic encryption (TFHE) [BGG+18]. We also give a computationally more efficient protocol based on weaker assumptions, namely n -out-of- n threshold additively homomorphic encryption (TAHE) [Ben94, Pai99], with communication complexity of $\tilde{O}(nT)$ that almost matches the lower bound.¹ All these protocols achieve semi-honest security where up to $(n - 1)$ parties could be corrupted.

Table 1. Communication lower and upper bounds for multi-party threshold PSI.

Functionality	Communication lower bound	TFHE-based upper bound	TAHE-based upper bound
$\mathcal{F}_{\text{TPSI-int}}$	$\Omega(nT)$	$O(nT)$	Unknown
$\mathcal{F}_{\text{TPSI-diff}}$	$\Omega(nT)$	$O(nT)$	$\tilde{O}(nT)$

Our Protocols. As summarized in Table 1, we present three protocols for upper bounds, one for $\mathcal{F}_{\text{TPSI-int}}$ and two for $\mathcal{F}_{\text{TPSI-diff}}$. At a high level, all three protocols compute their functionality in two phases. In the first phase, they perform

¹ $\tilde{O}(\cdot)$ hides polylog factors. All the upper bounds omit a $\text{poly}(\lambda)$ factor where λ is the security parameter.

a multi-party private intersection cardinality testing where the parties jointly decide whether their intersection is sufficiently large. In particular, for $\mathcal{F}_{\text{TPSI-int}}$, the cardinality testing, which we call $\mathcal{F}_{\text{CTest-int}}$, allows all the parties to learn whether $|I| \geq (m - T)$. For $\mathcal{F}_{\text{TPSI-diff}}$, the cardinality testing, which we call $\mathcal{F}_{\text{CTest-diff}}$, allows all the parties to learn whether $|\bigcup_{i=1}^n S_i \setminus I| \leq T$. The communication complexity of our protocols for $\mathcal{F}_{\text{CTest-int}}$ and $\mathcal{F}_{\text{CTest-diff}}$ is summarized in Table 2. In particular, for $\mathcal{F}_{\text{CTest-int}}$, we present a protocol with communication complexity $O(nT)$ based on TFHE. For $\mathcal{F}_{\text{CTest-diff}}$, we show a TFHE-based construction with communication complexity $O(nT)$ and a TAHE-based construction with communication complexity $\tilde{O}(nT)$.

Table 2. Communication complexity of our protocols for multi-party private cardinality testing.

Functionality protocol	TFHE-based protocol	TAHE-based
$\mathcal{F}_{\text{CTest-int}}$	$O(nT)$	Unknown
$\mathcal{F}_{\text{CTest-diff}}$	$O(nT)$	$\tilde{O}(nT)$

If the intersection is sufficiently large, namely it passes the cardinality testing, then the parties start the second phase of our protocols, which allows each party P_i to learn their set difference $S_i \setminus I$. We present a single protocol for the second phase, which works for both $\mathcal{F}_{\text{TPSI-int}}$ and $\mathcal{F}_{\text{TPSI-diff}}$. The second-phase protocol is based on TAHE and has communication complexity of $O(nT)$. Thus, to construct a protocol for multi-party threshold PSI, we combine the first-phase protocols summarized in Table 2 with the second-phase one described above. Doing so, we achieve the communication upper bounds in Table 1.

This modular design enables our constructions to minimize the use of TFHE as it is not needed in the second phase. Moreover, it allows future work to focus on improving Table 2. In particular, to design a protocol for $\mathcal{F}_{\text{TPSI-int}}$ from assumptions weaker than TFHE, future work could focus on building protocols for $\mathcal{F}_{\text{CTest-int}}$ and directly plug in our second phase protocol after that.

Communication Topology. All our protocols are designed in the so-called *star network* topology, where a designated party communicates with every other party. An added benefit of this topology is that not all parties must be online at the same time. Our communication lower bounds are proved in *point-to-point* fully connected networks, which are a generalization of the star network.

For networks with broadcast channels, we prove another communication lower bound of $\Omega(T \log n + n)$ for $\mathcal{F}_{\text{TPSI-int}}$ in the full version and leave further exploration in the broadcast model for future work.

1.2 Other Implications

Two-Party Threshold PSI. Recall that in the two-party case, both functionalities $\mathcal{F}_{\text{TPSI-int}}$ and $\mathcal{F}_{\text{TPSI-diff}}$ are identical. Ghosh and Simkin [GS19a] built

a two-party threshold PSI protocol from AHE with communication complexity $\tilde{O}(T^2)$. They left it as an open problem to build a two-party threshold PSI protocol with communication complexity $\tilde{O}(T)$ from assumptions weaker than FHE. Observe that for the special case of $n = 2$, we can achieve a two-party threshold PSI protocol with communication complexity $\tilde{O}(T)$ from AHE thereby solving this open problem (refer to Sect. 6 and Sect. 7 for more details).

Sublinear Communication PSI. Our multi-party threshold PSI protocols for both $\mathcal{F}_{\text{TPSI-int}}$ and $\mathcal{F}_{\text{TPSI-diff}}$ can also be used to achieve multi-party “regular” PSI² where the communication complexity only grows with the size of the set difference and independent of the input set sizes. In particular, if we run a sequence of multi-party threshold PSI protocols on $T = 2^0, 2^1, 2^2, \dots$ until hitting the smallest $T = 2^k$ where the protocol outputs the intersection, then we can achieve multi-party PSI. The communication complexity of the resulting protocol is a factor $\log T$ times that of a single instance but still independent of the input set sizes. Therefore, when the intersection is very large, namely the set difference is significantly smaller than the set sizes, this new approach achieves the *first* multi-party PSI with sublinear (in the set sizes) communication complexity.

Compact MPC. It is an open problem to construct a *compact* MPC protocol in the plain model where the communication complexity does not grow with the output length of the function. Prior works [HW15, BFK+19] construct compact MPC for general functions in the presence of a trusted setup (CRS, random oracle) from strong computational assumptions such as obfuscation. Our multi-party threshold PSI protocols have communication complexity independent of the output size (the set intersection). To the best of our knowledge, ours are the *first* compact MPC protocols for any non-trivial function in the plain model. The only prior compact protocol in the plain model we are aware of is the two-party threshold PSI protocol [GS19a].

1.3 Concurrent and Independent Work

Concurrent to our work, a recent update to the full version of the paper by Ghosh and Simkin [GS19b] extends the two-party threshold PSI protocol to the multi-party setting and consider the functionality $\mathcal{F}_{\text{TPSI-int}}$. They do not consider the functionality $\mathcal{F}_{\text{TPSI-diff}}$ that we additionally consider in our work. For $\mathcal{F}_{\text{TPSI-int}}$, [GS19b] also first constructs a TFHE-based protocol for the intersection cardinality testing $\mathcal{F}_{\text{CTest-int}}$ with communication complexity $O(nT)$. Then in the second phase for computing the intersection, they use an MPC protocol to compute the evaluations of a random polynomial, where the communication complexity depends on how the MPC is instantiated, which is not discussed

Another concurrent work by Branco, Döttling, and Pu [BDP21] studies multi-party private intersection cardinality testing with the functionality $\mathcal{F}_{\text{CTest-int}}$ and presents a TAHE-based protocol with communication complexity $\tilde{O}(nT^2)$, which

² By “regular” PSI, we refer to the standard notion of PSI without threshold.

complements our Table 2. They also do not consider the other functionality $\mathcal{F}_{\text{CTest-diff}}$.

1.4 Roadmap

We describe some notations and definitions in Sect. 2, a technical overview in Sect. 3, and the lower bound in Sect. 4. We present the TFHE based protocols for $\mathcal{F}_{\text{CTest-int}}$ and $\mathcal{F}_{\text{CTest-diff}}$ in Sect. 5 and the TAHE based protocol for $\mathcal{F}_{\text{CTest-diff}}$ in Sect. 6. We present the second phase protocol to compute the actual intersection in Sect. 7.

2 Preliminaries

In this section, we introduce some notations and define the our ideal functionalities. See the full version for the remaining definitions.

2.1 Notations

We use λ to denote the security parameters. By $\text{poly}(\lambda)$ we denote a polynomial function in λ . By $\text{negl}(\lambda)$ we denote a negligible function, that is, a function f such that $f(\lambda) < 1/p(\lambda)$ holds for any polynomial $p(\cdot)$ and sufficiently large λ . We use $\llbracket x \rrbracket$ to denote an encryption of x . We use $\tilde{O}(x)$ to ignore any polylog factor, namely $\tilde{O}(x) = O(x \cdot \text{polylog}(x))$.

2.2 Multi-party Threshold Private Set Intersection

Setting. Consider n parties P_1, \dots, P_n with input sets S_1, \dots, S_n respectively. Throughout the paper, we consider all the sets to be of equal size m . We assume that the set elements come from a field \mathbb{F}_p , where p is a $\Theta(\lambda)$ -bit prime. Also, throughout the paper, we focus only on the point-to-point network channels. For the lower bounds, we consider a setting where every pair of parties has a point-to-point channel between them. For the upper bounds, we consider a more restrictive model – the star network, where only one central party has a point-to-point channel with every other party and the other parties cannot communicate with each other.

The goal of the parties is to run an MPC protocol Π at the end of which each party learns the intersection I of all the sets if certain conditions hold. In the definition of two-party threshold PSI, both parties P_1 and P_2 learn the intersection I if the size of their set difference is small, namely $|(S_1 \setminus S_2) \cup (S_2 \setminus S_1)| < 2T$. In the multi-party case, we consider two different functionalities, each of which might be better suited to different applications.

Functionalities. In the first definition, we consider functionality $\mathcal{F}_{\text{TPSI-int}}$, in which each party P_i learns the intersection I if the size of its own set minus the intersection is small, namely $|S_i \setminus I| \leq T$ for some threshold T . Recall that we

consider all the sets to be of equal size, hence either all the parties learn the output or all of them don't. In the second definition, we consider a functionality $\mathcal{F}_{\text{TPSI-diff}}$, where each party learns the intersection I if the size of the union of all the sets minus the intersection is small, namely $|(\bigcup_{i=1}^n S_i) \setminus I| \leq T$. The formal definitions of the two ideal functionalities are shown in Fig. 1 and Fig. 2.

Parameters: Parties P_1, \dots, P_n . Each party has a set of m elements. Threshold $T \in \mathbb{N}$.
Inputs: Party P_i has an input set $S_i = \{a_1^i, \dots, a_m^i\}$ where $a_j^i \in \mathbb{F}_p$ for all $j \in [m]$.
Output: Each party P_i receives $I = \bigcap_{i=1}^n S_i$ if and only if $|S_i \setminus I| \leq T$.

Fig. 1. Ideal functionality $\mathcal{F}_{\text{TPSI-int}}$ for multi-party threshold PSI.

Parameters: Parties P_1, \dots, P_n . Each party has a set of m elements. Threshold $T \in \mathbb{N}$.
Inputs: Party P_i has an input set $S_i = \{a_1^i, \dots, a_m^i\}$ where $a_j^i \in \mathbb{F}_p$ for all $j \in [m]$.
Output: Each party P_i receives $I = \bigcap_{i=1}^n S_i$ if and only if $(\bigcup_{i=1}^n S_i \setminus I) \leq T$.

Fig. 2. Ideal functionality $\mathcal{F}_{\text{TPSI-diff}}$ for multi-party threshold PSI.

2.3 Multi-party Private Intersection Cardinality Testing

An important building block in our multi-party threshold PSI protocols is a multi-party protocol for private intersection cardinality testing which we define below. Consider n parties P_1, \dots, P_n with input sets S_1, \dots, S_n respectively of equal size m . Their goal is to run an MPC protocol Π at the end of which each party learns whether the size of the intersection I of all the sets is sufficiently large. As before, we consider two functionalities. In the first functionality $\mathcal{F}_{\text{CTest-int}}$, each party P_i learns whether $|S_i \setminus I| \leq T$. In the second functionality $\mathcal{F}_{\text{CTest-diff}}$, each party learns whether $|(\bigcup_{i=1}^n S_i) \setminus I| \leq T$. The formal definitions of the two ideal functionalities are presented in Fig. 3 and Fig. 4.

Parameters: Parties P_1, \dots, P_n . Each party has a set of m elements. Threshold $T \in \mathbb{N}$.
Inputs: Party P_i has an input set $S_i = \{a_1^i, \dots, a_m^i\}$ where $a_j^i \in \mathbb{F}_p$ for all $j \in [m]$.
Output: Each party P_i receives similar if $|S_i \setminus I| \leq T$ and different otherwise where $I = \bigcap_{i=1}^n S_i$.

Fig. 3. Ideal functionality $\mathcal{F}_{\text{CTest-int}}$ for multi-party intersection cardinality test.

Parameters: Parties P_1, \dots, P_n . Each party has a set of m elements. Threshold $T \in \mathbb{N}$.
Inputs: Party P_i has an input set $S_i = \{a_1^i, \dots, a_m^i\}$ where $a_j^i \in \mathbb{F}_p$ for all $j \in [m]$.
Output: Each party P_i receives similar if $|\left(\bigcup_{i=1}^n S_i\right) \setminus I| \leq T$ and different otherwise where $I = \bigcap_{i=1}^n S_i$.

Fig. 4. Ideal functionality $\mathcal{F}_{\text{CTest-diff}}$ for multi-party intersection cardinality test.

3 Technical Overview

We now give an overview of the techniques used in our work. We denote P_1 as the designated party that can communicate with all the other parties.

3.1 TFHE-Based Protocol for $\mathcal{F}_{\text{CTest-int}}$

In Sect. 5.1 we construct a protocol for $\mathcal{F}_{\text{CTest-int}}$ from TFHE. Our starting point is the two-party protocol of [GS19a]. Recall that there are two parties Alice and Bob with sets $S_A = \{a_1, \dots, a_m\}$ and $S_B = \{b_1, \dots, b_m\}$ respectively. These sets define two polynomials $p_A(x) := \prod_{i=1}^m (x - a_i)$ and $p_B(x) := \prod_{i=1}^m (x - b_i)$. Let $I := S_A \cap S_B$ be the intersection. A key observation in [MTZ03, GS19a] is that $p(x) := \frac{p_B(x)}{p_A(x)} = \frac{p_{B \setminus I}(x)}{p_{A \setminus I}(x)}$. Both the numerator and denominator of p have degree $m - |I|$. If $m - |I| = |S_A \setminus I| \leq T$, then $p(x)$ has degree at most $2T$ and can be recovered from $2T + 1$ evaluations by rational function interpolation.³ Given $p(x)$, the elements in $S_A \setminus I$ are simply the roots of the polynomial in the denominator.

Two-Party Protocol. At a high level, the two-party protocol [GS19a] works as follows. First, Alice and Bob evaluate their own polynomials on $2T + 1$ publicly known distinct points $\{\alpha_1, \dots, \alpha_{2T+1}\}$ to obtain $\{p_A(\alpha_1), \dots, p_A(\alpha_{2T+1})\}$ and $\{p_B(\alpha_1), \dots, p_B(\alpha_{2T+1})\}$, respectively. Then, Alice generates a public-secret key pair for FHE and sends Bob the FHE public key, encrypted evaluations $\{\llbracket p_A(\alpha_1) \rrbracket, \dots, \llbracket p_A(\alpha_{2T+1}) \rrbracket\}$, a uniformly random z and encrypted evaluation $\llbracket p_A(z) \rrbracket$. Bob can homomorphically interpolate the rational function $\llbracket p(x) \rrbracket$ from $\{\llbracket p_A(\alpha_1) \rrbracket, \dots, \llbracket p_A(\alpha_{2T+1}) \rrbracket\}$ and $\{p_B(\alpha_1), \dots, p_B(\alpha_{2T+1})\}$, and then homomorphically compute $\llbracket p(z) \rrbracket$. Bob can also compute $p_B(z)$ and homomorphically compute $\frac{p_B(z)}{\llbracket p_A(z) \rrbracket}$. We know that $p(z) = \frac{p_B(z)}{p_A(z)}$ if and only if the degree of $p(x)$ is $\leq 2T$. Therefore Bob homomorphically computes an encryption of the predicate $\llbracket b \rrbracket := \left(\llbracket p(z) \rrbracket \stackrel{?}{=} \frac{p_B(z)}{\llbracket p_A(z) \rrbracket} \right)$ and sends the encryption $\llbracket b \rrbracket$ back to Alice. Finally Alice decrypts and learns b .

³ A rational function is a fraction of two polynomials. We refer to Minsky et al. [MTZ03] for details on rational function interpolation over a field. Also, we note that monic polynomials can be interpolated with $2T$ evaluation but we use $2T + 1$ for consistency with our other protocols.

Multi-party Protocol. For n parties, a natural idea is to consider

$$p(x) := \frac{p_2(x) + \dots + p_n(x)}{p_1(x)} = \frac{p_{2 \setminus I}(x) + \dots + p_{n \setminus I}(x)}{p_{1 \setminus I}(x)}, \tag{1}$$

where $p_i(x)$ encodes the set $S_i = \{a_1^i, \dots, a_m^i\}$ as $p_i(x) := \prod_{j=1}^m (x - a_j^i)$. The n parties first jointly generate the TFHE keys. Each party P_i sends encrypted evaluations $\{\llbracket p_i(\alpha_1) \rrbracket, \dots, \llbracket p_i(\alpha_{2T+1}) \rrbracket, \llbracket p_i(z) \rrbracket\}$ to P_1 . Now P_1 can interpolate $\llbracket p(x) \rrbracket$ from $2T + 1$ evaluations and compute an encryption $\llbracket b \rrbracket := \left(\llbracket p(z) \rrbracket \stackrel{?}{=} \frac{\llbracket p_2(z) \rrbracket + \dots + \llbracket p_n(z) \rrbracket}{p_1(z)} \right)$. Finally the parties jointly decrypt $\llbracket b \rrbracket$.

Unexpected Degree Reduction. This seemingly correct protocol has a subtle issue.⁴ Intuitively, we want to argue that $p(x)$ in Eq. 1 has degree $\leq 2T$ if and only if $|S_1 \setminus I| \leq T$. However, this is not true because elements not in the intersection might be accidentally canceled out, which results in a lower degree than the intersection cardinality would imply. As a concrete example, consider three sets with distinct elements $S_1 = \{a\}$, $S_2 = \{b\}$, $S_3 = \{c\}$, where $b + c = 2 \cdot a$. The intersection $I = \emptyset$. Ideally we hope the rational polynomial $p(x)$ has degree 1 in both the numerator and denominator because $|S_1 \setminus I| = 1$. However,

$$p(x) = \frac{(x - b) + (x - c)}{x - a} = \frac{2x - (b + c)}{x - a} = \frac{2x - 2a}{x - a} = 2.$$

Randomness to the Rescue. On first thought, this approach seems fundamentally flawed as additional roots can always be created if we add polynomials in the numerator. To solve this problem, we add a random multiplicative term $(x - r_i)$ to each polynomial p_i and set a new polynomial $p'_i(x) := p_i(x) \cdot (x - r_i)$ for a random r_i chosen by party P_i . Now, consider the rational polynomial

$$p'(x) := \frac{p'_2(x) + \dots + p'_n(x)}{p'_1(x)} = \frac{p'_{2 \setminus I}(x) + \dots + p'_{n \setminus I}(x)}{p'_{1 \setminus I}(x)}.$$

At a high level, the terms $(x - r_i)$ will randomize the roots of the numerator sufficiently to ensure that these roots are unlikely to coincide with the roots of the denominator.

3.2 TFHE-Based Protocol for $\mathcal{F}_{\text{CTest-diff}}$

In Sect. 5.2 we present an TFHE-based protocol for $\mathcal{F}_{\text{CTest-diff}}$. In summary, party P_1 tries to homomorphically interpolate

$$\tilde{p}_i(x) = \frac{p_i(x)}{p_1(x)} = \frac{p_{i \setminus 1}(x)}{p_{1 \setminus i}(x)}$$

⁴ In fact, this subtle issue was initially overlooked by [GS19b] in their recent update of the multi-party protocol. It has subsequently been fixed after we notified them.

from $(2T+1)$ evaluations and computes encrypted $D_{1,i} = S_1 \setminus S_i$ as well as $D_{i,1} = S_i \setminus S_1$ for every other party P_i . Note that if $|\bigcup_{i=1}^m S_i \setminus I| \leq T$, then $|S_i \setminus I| \leq T$ for all i and the degree of each $\tilde{p}_i(x)$ is at most $2T$, hence P_1 can interpolate it using $(2T+1)$ evaluations. Observe that $(\bigcup_{i=1}^m S_i) \setminus I = \bigcup_{i=2}^m (D_{1,i} \cup D_{i,1})$, because each element $a \in (\bigcup_{i=1}^m S_i) \setminus I$ must be one of the two cases: (1) $a \in S_1$ and $a \notin S_i$ for some i (i.e., $a \in D_{1,i}$), or (2) $a \notin S_1$ and $a \in S_i$ for some i (i.e., $a \in D_{i,1}$). Therefore, party P_1 can homomorphically compute an encryption of $(\bigcup_{i=1}^m S_i) \setminus I$ and an encryption of the predicate $b = \left(\left| \left(\bigcup_{i=1}^m S_i \right) \setminus I \right| \stackrel{?}{\leq} T \right)$. Finally, as before, the n parties jointly decrypt $\llbracket b \rrbracket$ to learn the output.

3.3 TAHE-Based Protocol for $\mathcal{F}_{\text{CTest-diff}}$

Section 6 presents our protocol for $\mathcal{F}_{\text{CTest-diff}}$ based on TAHE. This protocol reduces the communication complexity for two-party from $\tilde{O}(T^2)$ to $\tilde{O}(T)$ as well as generalizes it to multi-party with communication $\tilde{O}(Tn)$.

Two-Party Protocol. For two parties Alice and Bob with private sets S_A and S_B , if we encode their elements into two polynomials $p_A(x) = \sum_{i=1}^m x^{a_i}$ and $p_B(x) = \sum_{i=1}^m x^{b_i}$, then the number of monomials in the polynomial $p(x) := p_A(x) - p_B(x)$ is exactly $|(S_A \setminus S_B) \cup (S_B \setminus S_A)|$. Now the problem of cardinality testing (i.e., determining if $|(S_A \setminus S_B) \cup (S_B \setminus S_A)| \leq 2T$) has been reduced to determining whether the number of monomials in $p(x)$ is $\leq 2T$. Using the polynomial sparsity test of Grigorescu et al. [GJR10], we can further reduce the problem to determining whether the Hankel matrix below is singular or not:

$$H = \begin{bmatrix} p(u^0) & p(u^1) & \dots & p(u^{2T}) \\ p(u^1) & p(u^2) & \dots & p(u^{2T+1}) \\ \vdots & \vdots & \ddots & \vdots \\ p(u^{2T}) & p(u^{2T+1}) & \dots & p(u^{4T}) \end{bmatrix},$$

where u is chosen uniformly at random. In the two-party protocol, Alice generates a public-secret key pair for AHE and sends Bob the public key, a uniformly random u along with encrypted Hankel matrix for p_A . Then Bob can homomorphically compute encrypted Hankel matrix for p . Now Alice holds the secret key and Bob holds an encryption of matrix H . They need to jointly perform a secure matrix singularity testing to determine if the matrix is singular, which can be done using the protocol of Kiltz et al. [KMWF07] with communication $\tilde{O}(T^2)$.

Our Approach. Our key observation is that the protocol of Kiltz et al. [KMWF07] can be used to perform singularity testing for *arbitrary* matrices, while we are only interested in testing the singularity of Hankel matrices. Since a Hankel matrix only has linear (in its dimension) number of distinct entries, there is a more efficient way to test its singularity. In particular, the work of Brent et al. [BGY80] demonstrates an elegant connection between the problem of testing singularity of a Hankel matrix and the so-called “half-GCD” problem, which can be solved in quasi-linear time. Thus, testing singularity of

the Hankel matrix H only takes $\tilde{O}(T)$ computation. In our scenario, we can first let Alice and Bob learn an additive share of H , and then engage in a two-party computation (using AHE or Yao’s garbled circuits) to jointly test if H is singular or not. The important point to note here is that both communication and computation are only quasi-linear in the dimension of H . This is already an improvement over the quadratic cost of protocol in [KMWF07] and solves the open problem posed by Ghosh and Simkin [GS19a].

Multi-party Protocol. In designing a multi-party protocol, our strategy is to first find a polynomial where the number of monomials equals the size of the set difference $|(\bigcup_{i=1}^m S_i) \setminus I|$. Furthermore, the polynomial should only involve linear operations among the parties, which allows the parties to obtain additive secret shares of the Hankel matrix for the polynomial. Then, the parties perform an MPC protocol to test singularity of the Hankel matrix.

3.4 Computing Set Intersection

In Sect. 7 we present a single construction that computes the concrete set intersection for both $\mathcal{F}_{\text{TPSI-int}}$ and $\mathcal{F}_{\text{TPSI-diff}}$ after the cardinality testing.

Two-Party Protocol. For two parties Alice and Bob, we use the first encoding method to encode the elements into two polynomials $p_A(x) = \prod_{i=1}^m (x - a_i)$ and $p_B(x) = \prod_{i=1}^m (x - b_i)$. After the cardinality testing, we already know that the rational polynomial $p(x) := \frac{p_B(x)}{p_A(x)} = \frac{p_{B \setminus I}(x)}{p_{A \setminus I}(x)}$ has degree at most $2T$. If Alice learns the evaluation of $p_B(\cdot)$ on $2T + 1$ distinct points $\{\alpha_1, \dots, \alpha_{2T+1}\}$, then she can evaluate p_A on those points by herself and compute $\{p(\alpha_1), \dots, p(\alpha_{2T+1})\}$. Using these evaluations of $p(\cdot)$, Alice can recover $p(x)$ by rational polynomial interpolation, and then learn the set difference $S_A \setminus I$ from the denominator of $p(x)$. However, $p(x)$ also allows Alice to learn $S_B \setminus I$, which breaks security. Instead of letting Alice learn the evaluations of $p_B(\cdot)$, the two-party protocol of [GS19a] enables Alice to learn the evaluations of a “noisy” polynomial $V(x) := p_A(x) \cdot R_1(x) + p_B(x) \cdot R_2(x)$, where R_1 and R_2 are uniformly random polynomials of degree T . Note that

$$p'(x) := \frac{V(x)}{p_A(x)} = \frac{p_{A \setminus I}(x) \cdot R_1(x) + p_{B \setminus I}(x) \cdot R_2(x)}{p_{A \setminus I}(x)}$$

has degree at most $3T$. Given $3T + 1$ evaluations of $V(\cdot)$, Alice can interpolate $p'(x)$ and figure out the denominator, but now the numerator is sufficiently random and does not leak any other information about S_B .

Multi-party Protocol. For n parties, we first encode each set $S_i = \{a_1^i, \dots, a_m^i\}$ as a polynomial $p_i(x) := \prod_{j=1}^m (x - a_j^i)$, and then define

$$\begin{aligned} V(x) &:= p_1(x) \cdot R_1(x) + \dots + p_n(x) \cdot R_n(x) \\ &:= p_1(x) \cdot (R_{1,1}(x) + \dots + R_{n,1}(x)) + \dots + p_n(x) \cdot (R_{1,n}(x) + \dots + R_{n,n}(x)), \end{aligned}$$

where $(R_{i,1}, \dots, R_{i,n})$ are random polynomials of degree T generated by party P_i . Different from the two-party protocol, it is crucial that each party P_i contributes a random term in every polynomial R_1, \dots, R_n . For both functionalities $\mathcal{F}_{\text{TPSI-int}}$ and $\mathcal{F}_{\text{TPSI-diff}}$, if the protocol passes the cardinality testing, then

$$p'(x) := \frac{V(x)}{p_1(x)} = \frac{p_{1 \setminus I}(x) \cdot R_1(x) + \dots + p_{n \setminus I}(x) \cdot R_n(x)}{p_{1 \setminus I}(x)}$$

has degree at most $3T$. If P_1 learns $3T + 1$ evaluations of $V(\cdot)$, then it can interpolate $p'(x)$ and recover $S_1 \setminus I$ from the denominator while the numerator does not leak any other information. Since $V(\cdot)$ can be broken down to linear operations among the parties, it can be securely evaluated by TAHE.

Communication Blow-Up. However, this protocol requires $O(n^2)$ communication complexity per evaluation, and the total communication complexity is $O(n^2T)$ for $(3T + 1)$ evaluations. Observe that the bottleneck of the communication in this approach is that every party P_i needs to contribute n randomizing polynomials $(R_{i,1}, \dots, R_{i,n})$. Through a careful analysis we demonstrate that it is sufficient for each party to only contribute two randomizing polynomials. The first is used to randomize their own polynomial while the second randomizes the polynomials from the other parties. Nevertheless, there is a subtle issue of unexpected degree reduction, similar to what we have seen in the TFHE-based protocol $\mathcal{F}_{\text{CTest-int}}$. We follow the same approach as in the TFHE-based protocol by adding additional randomness in the polynomial, which reduces the communication complexity to $O(nT)$.

3.5 Lower Bounds

We briefly discuss the communication lower bound for multi-party threshold PSI. To prove lower bound in the point-to-point network, we perform a reduction from two-party threshold PSI (for which [GS19a] showed a lower bound of $\Omega(T)$) to multi-party threshold PSI. We first prove that the total “communication complexity of any party” is $\Omega(T)$ which denotes the sum of all the bits exchanged by that party (both sent and received). As a corollary, the total communication complexity of any multi-party threshold PSI protocol is $\Omega(nT)$. We refer to Sect. 4 for more details about the reduction.

To prove a lower bound in the broadcast model, we rely on the communication lower bound of the multi-party set disjointness problem shown by Braverman and Oshman [BO15]. We reduce the problem of multi-party set disjointness to multi-party threshold PSI $\mathcal{F}_{\text{TPSI-int}}$ and prove a lower bound $\Omega(T \log n + n)$ for any multi-party threshold PSI protocol in the broadcast network. We refer to the full version for more details about the reduction.

4 Communication Lower Bound

In this section, we prove communication lower bounds for multi-party threshold PSI protocols in the point-to-point network model. Recall that we consider all

parties to have sets of the same size m . We show that any secure protocol must have communication complexity at least $\Omega(n \cdot T)$ for both functionalities $\mathcal{F}_{\text{TPSI-int}}$ and $\mathcal{F}_{\text{TPSI-diff}}$. We prove the lower bound for $\mathcal{F}_{\text{TPSI-int}}$ and defer the proof for $\mathcal{F}_{\text{TPSI-diff}}$ to the full version. Before proving the lower bound, we first prove another related theorem below.

Theorem 1. *For any multi-party threshold PSI protocol for functionality $\mathcal{F}_{\text{TPSI-int}}$ that is secure against a semi-honest adversary that can corrupt up to $(n-1)$ parties, for every party P_i , the communication complexity of P_i is $\Omega(T)$.⁵*

Proof. Suppose this is not true. That is, suppose there exists a secure multi-party threshold PSI protocol Π for functionality $\mathcal{F}_{\text{TPSI-int}}$ in which for some party P_{i^*} , $\text{CC}(P_{i^*}) = o(T)$ where $\text{CC}(\cdot)$ denotes the communication complexity. We will now use this protocol Π as a subroutine to design a secure two-party threshold PSI protocol which has communication complexity $o(T)$.

Consider two parties Q_1 and Q_2 with input sets X_1 and X_2 (of same size m) who wish to run a secure two-party threshold PSI protocol for the following functionality: both parties learn the output if $|(X_1 \setminus X_2) \cup (X_2 \setminus X_1)| \leq 2 \cdot T$. We invoke the multi-party threshold PSI protocol Π with threshold T as follows: Q_1 emulates the role of party P_{i^*} with input set $S_{i^*} = X_1$ and Q_2 emulates the role of all the other $(n-1)$ parties with each of their input sets as X_2 . From the definition of the functionality $\mathcal{F}_{\text{TPSI-int}}$, Q_1 learns the output at the end of the protocol if and only if $|X_1 \setminus I| \leq T$. Similarly, Q_2 learns the output at the end of the protocol if and only if $|X_2 \setminus I| \leq T$. Notice that since $|X_1| = |X_2|$ and $I = X_1 \cap X_2$, $|X_1 \setminus I| = |X_2 \setminus I|$. Thus, the parties learn the output if and only if $(|X_1 \setminus I|) + (|X_2 \setminus I|) \leq 2 \cdot T$, namely $|(X_1 \setminus X_2) \cup (X_2 \setminus X_1)| \leq 2 \cdot T$, which is the functionality of the two-party threshold PSI. Therefore, correctness is easy to observe. For security, notice that if Q_1 is corrupt, we can simulate it by considering only a corrupt P_{i^*} in the underlying protocol Π and if Q_2 is corrupt, we can simulate it by considering all parties except P_{i^*} to be corrupt in the underlying protocol Π .

Finally, notice that the communication complexity of the two-party protocol is exactly the same as $\text{CC}(P_{i^*})$ in the multi-party protocol Π , which is $o(T)$. However, recall from the work of Ghosh and Simkin [GS19a] that any two-party threshold PSI for this functionality has communication complexity lower bound $\Omega(T)$ leading to a contradiction. Thus, the assumption that there exists a secure multi-party PSI protocol Π in which for some party P_{i^*} , $\text{CC}(P_{i^*}) = o(T)$ is wrong and this completes the proof of the theorem.

It is easy to observe that as a corollary of the above theorem, in a setting with only point-to-point channels (which also includes the star network), the overall communication complexity of the protocol must be at least n times the minimum communication complexity that each party is involved in, giving the lower bound of $\Omega(n \cdot T)$. Formally,

⁵ We define the communication complexity of a party P_i in any protocol execution as the complexity of all the communication that P_i is involved in. That is, the complexity of the messages both incoming to and outgoing from P_i .

Corollary 1. *For any multi-party threshold PSI protocol for functionality $\mathcal{F}_{\text{TPSI-int}}$ that is secure against a semi-honest adversary that can corrupt up to $(n - 1)$ parties, the communication complexity is $\Omega(n \cdot T)$.*

5 TFHE-Based Private Intersection Cardinality Testing

In this section, we present two protocols for private intersection cardinality testing, one for functionalities $\mathcal{F}_{\text{CTest-int}}$ (described in Fig. 3) and the other for $\mathcal{F}_{\text{CTest-diff}}$ (described in Fig. 4). Both protocols are based on n -out-of- n threshold fully homomorphic encryption with distributed setup. The former functionality states that the intersection must be of size at least $(m - T)$ where m is the size of each set. The latter functionality requires the difference between the union of all the sets and the intersection be of size at most T . Due to the possibility of elements appearing in a strict subset of the sets, these two functionalities are not equivalent.

5.1 Protocol for Functionality $\mathcal{F}_{\text{CTest-int}}$

In this protocol, we compute the cardinality predicate b where $b = 1$ if and only if $\forall i, |S_i \setminus I| \leq T$. The communication complexity of this protocol involves sending $O(nT)$ TFHE ciphertexts and performing a single decryption of the result. We briefly describe the approach below.

Each party P_i first encodes their set S_i as a polynomial $\mathbf{p}_i(x) := \prod_{a \in S_i} (x - a) \in \mathbb{F}[x]$. Each of these polynomials are then randomized as $\mathbf{p}'_i(x) := \mathbf{p}_i(x) \cdot (x - r_i)$ where P_i uniformly samples $r_i \xleftarrow{\$} \mathbb{F}$. The central party also picks a random $z \xleftarrow{\$} \mathbb{F}$ which is sent to every other party. Each party P_i then computes $e_{i,j} := \mathbf{p}'_i(j)$ for $j \in [2T + 3]$ and $e'_i := \mathbf{p}'_i(z)$. P_i sends the ciphertexts $\llbracket e_{i,j} \rrbracket := \text{TFHE.Enc}(\text{pk}, e_{i,j})$ and $\llbracket e'_i \rrbracket := \text{TFHE.Enc}(\text{pk}, e'_i)$ to P_1 . Party P_1 considers the rational polynomial

$$\mathbf{p}'(x) = \frac{\mathbf{p}'_2(x) + \dots + \mathbf{p}'_n(x)}{\mathbf{p}'_1(x)}$$

and homomorphically computes $2T + 3$ encrypted evaluations

$$\left(j, \llbracket \frac{e_{2,j} + \dots + e_{n,j}}{e_{1,j}} \rrbracket \right)$$

for $j \in [2T + 3]$. Using these encrypted evaluations, P_1 homomorphically computes an encrypted rational polynomial $\llbracket \mathbf{p}^*(x) \rrbracket$ using rational polynomial interpolation. Note that $\mathbf{p}^*(x) = \mathbf{p}'(x)$ if $\mathbf{p}'(x)$ has degree at most $2T + 2$. Furthermore, P_1 can homomorphically compute an encryption of the predicate $b := \left(\mathbf{p}^*(z) \stackrel{?}{=} \frac{e'_2 + \dots + e'_n}{e'_1} \right)$. Finally the parties jointly perform a threshold decryption of $\llbracket b \rrbracket$ and party P_1 learns the output which is sent to every other party. The full protocol is detailed in Fig. 5.

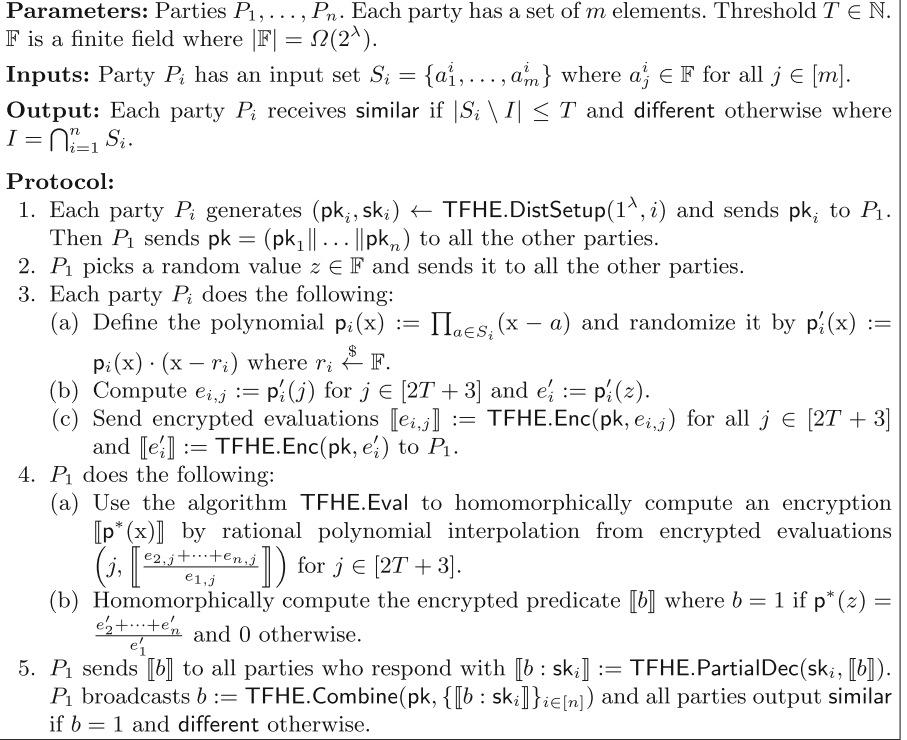


Fig. 5. Multi-party private intersection cardinality testing protocol $\Pi_{\text{TFHE-CTest-int}}$ for $\mathcal{F}_{\text{CTest-int}}$.

Theorem 2. *Assuming threshold FHE with distributed setup, protocol $\Pi_{\text{TFHE-CTest-int}}$ (Fig. 5) securely realizes $\mathcal{F}_{\text{CTest-int}}$ (Fig. 3).*

Proof. Correctness. We first prove the protocol is correct. By the correctness of the TFHE scheme, we only need to show that the computed predicate $b = 1$ if and only if $\forall i, |S_i \setminus I| \leq T$. First consider the case where the protocol should output similar. Since

$$\mathbf{p}'(x) = \frac{\mathbf{p}'_2(x) + \dots + \mathbf{p}'_n(x)}{\mathbf{p}'_1(x)} = \frac{\mathbf{p}_{2 \setminus I}(x) \cdot (x - r_2) + \dots + \mathbf{p}_{n \setminus I}(x) \cdot (x - r_n)}{\mathbf{p}_{1 \setminus I}(x) \cdot (x - r_1)},$$

the degree of each term $\mathbf{p}_{i \setminus I}(x) \cdot (x - r_i)$ is at most $T + 1$ and therefore the rational polynomial interpolation requires a total of $(2T + 3)$ evaluation points. Therefore $\mathbf{p}^*(x) = \mathbf{p}'(x)$ and $\mathbf{p}^*(z) = \mathbf{p}'(z) = \frac{e'_2 + \dots + e'_n}{e'_1}$. Thus $b = 1$ as required.

Now consider the case where the protocol should output different, namely when $|I| < m - T$. Observe that $\gcd(\mathbf{p}_{1 \setminus I}, \dots, \mathbf{p}_{n \setminus I}) = 1$ by construction and therefore

$$\gcd(\mathbf{p}'_{2 \setminus I}(x) + \dots + \mathbf{p}'_{n \setminus I}(x), \mathbf{p}'_{1 \setminus I}(x)) = 1$$

except with negligible probability, where $\mathbf{p}'_{i \setminus I}(x) := \mathbf{p}_{i \setminus I}(x) \cdot (x - r_i)$. The algebraic proof is deferred to the full version.

Assuming $\gcd(\mathbf{p}'_{2 \setminus I}(x) + \dots + \mathbf{p}'_{n \setminus I}(x), \mathbf{p}'_{1 \setminus I}(x)) = 1$, it then follows that the degree of the rational polynomial $\mathbf{p}'(x)$ is the degree of $\mathbf{p}'_{2 \setminus I}(x) + \dots + \mathbf{p}'_{n \setminus I}(x)$ plus the degree of $\mathbf{p}'_{1 \setminus I}(x)$. The former must have a leading term with degree $(m - |I| + 1) > (T + 1)$. Similarly, the latter also has degree $(m - |I| + 1) > T + 1$. Hence the degree of $\mathbf{p}'(x)$ is at least $2T + 4$. The probability of $b = 1$ is $\Pr_z[\mathbf{p}'(z) = \mathbf{p}^*(z)]$ where $\mathbf{p}^*(x)$ is the polynomial interpolated by P_1 using $(2T + 3)$ evaluations. However, since the degree of $\mathbf{p}'(x)$ is at least $2T + 4$, $\Pr_z[\mathbf{p}'(z) = \mathbf{p}^*(z)] \leq \text{negl}(\lambda)$.

Communication Cost. Each party sends $(2T + 4)$ TFHE encryptions and one partial decryption to P_1 where each plaintext is a field element. P_1 sends one ciphertext to every other party. The size of each encryption and each partial decryption is $\text{poly}(\lambda)$. Thus, the overall communication complexity is $O(n \cdot T \cdot \text{poly}(\lambda))$ in a star network and the protocols runs in $O(1)$ rounds.

Security. Consider an environment \mathcal{Z} who corrupts a set \mathcal{S}^* of n^* parties where $n^* < n$. The simulator Sim has output $w \in \{\text{similar}, \text{different}\}$ from the ideal functionality. Sim sets a bit $b^* = 1$ if $w = \text{similar}$ and $b^* = 0$ otherwise. Also, for each corrupt party P_i , Sim has as input the tuple (S_i, r_i) indicating the party's input and randomness for the protocol. The strategy of the simulator Sim for our protocol is described below.

1. Sim runs the distributed key generation algorithm $\text{TFHE.DistSetup}(1^\lambda, i)$ of the TFHE scheme honestly on behalf of each honest party P_i as in the real world. Note that Sim also knows $(\{\text{sk}_i\}_{i \in \mathcal{S}^*})$ as it knows the randomness for the corrupt parties.
2. In Steps 2–4 of the protocol, Sim plays the role of the honest parties exactly as in the real world except that on behalf of every honest party P_i , whenever P_i has to send any ciphertext, compute $\llbracket 0 \rrbracket = \text{TFHE.Enc}(0)$ using fresh randomness.
3. In Step 5, on behalf of each honest party P_i , instead of sending the value $\llbracket b : \text{sk}_i \rrbracket$ by running the honest TFHE.PartialDec algorithm as in the real world, Sim computes the partial decryptions by running the simulator TFHE.Sim as follows: $\{\llbracket b : \text{Sim}_i \rrbracket\}_{i \in [n] \setminus \mathcal{S}^*} \leftarrow \text{TFHE.Sim}(\mathbf{C}, b^*, \llbracket b \rrbracket, \{\text{sk}_i\}_{i \in \mathcal{S}^*})$ where the circuit \mathbf{C} denotes the whole computation done by P_1 in the real world to evaluate bit b . On behalf of the honest party P_i the simulator sends $\llbracket b : \text{Sim}_i \rrbracket$. This corresponds to the ideal world.

We now show that the above simulation strategy is successful against all environments \mathcal{Z} that corrupt parties in a semi-honest manner. We will show this via a series of computationally indistinguishable hybrids where the first hybrid Hybrid_0 corresponds to the real world and the last hybrid Hybrid_2 corresponds to the ideal world.

- **Hybrid₀ - Real World:** In this hybrid, consider a simulator SimHyb that plays the role of the honest parties as in the real world.

- **Hybrid₁ - Simulate Partial Decryptions:** - In this hybrid, in Step 5, SimHyb simulates the partial decryptions generated by the honest parties as done in the ideal world. That is, the simulator calls $\{\llbracket b : \text{Sim}_i \rrbracket\}_{i \in [n] \setminus \mathcal{S}} \leftarrow \text{TFHE.Sim}(\mathbb{C}, b^*, \llbracket b \rrbracket, \{\text{sk}_i\}_{i \in \mathcal{S}})$. On behalf of the honest party P_i the simulator sends $\llbracket b : \text{Sim}_i \rrbracket$ instead of $\llbracket b : \text{sk}_i \rrbracket$.
- **Hybrid₂ - Switch Encryptions:** In this hybrid, SimHyb now computes every ciphertext generated on behalf of any honest party as encryptions of 0 as done by Sim in the ideal world. This hybrid corresponds to the ideal world.

We show that every pair of consecutive hybrids is computationally indistinguishable in the full version.

5.2 Protocol for Functionality $\mathcal{F}_{\text{CTest-diff}}$

This protocol will compute the cardinality predicate b where $b = 1$ if and only if $|(\bigcup_{i=1}^n S_i) \setminus I| \leq T$. The core idea behind the protocol is that P_1 (the star of the network) and P_i first run a protocol to compute an encryption (via TFHE) of their set differences $D_{1,i} = S_1 \setminus S_i$ and $D_{i,1} = S_i \setminus S_1$ with $O(T)$ communication complexity if $|S_1 \setminus S_i| \leq T$. Before we describe how this is achieved, notice that at this point, the protocol enables P_1 to reconstruct an encryption of $(\bigcup_{i=1}^n S_i) \setminus I = \bigcup_{i \in [n] \setminus \{1\}} (D_{1,i}^* \cup D_{i,1}^*)$ and a predicate b where $b = 1$ if and only if $|(\bigcup_{i=1}^n S_i) \setminus I| \leq T$. P_1 can then send this encryption to all parties to run threshold decryption.

We now describe in more detail how the encryption of $D_{1,i}$ and $D_{i,1}$ are computed. The idea follows from the two-party protocol of Ghosh and Simkin [GS19a]. Each party P_i encodes their set S_i as $\mathbf{p}_i(x) := \prod_{a \in S_i} (x - a) \in \mathbb{F}[x]$. P_i then computes $e_{i,j} := \mathbf{p}_i(j)$ for $j \in [2T + 1]$ and $e'_i := \mathbf{p}_i(z)$ on a special random point $z \in \mathbb{F}$ (picked uniformly at random by P_1). Party P_i encrypts these values as $\llbracket e_{i,j} \rrbracket, \llbracket e'_i \rrbracket$ and sends them to P_1 . Party P_1 considers the rational polynomial

$$\tilde{\mathbf{p}}_i(x) = \frac{\mathbf{p}_i(x)}{\mathbf{p}_1(x)} = \frac{\mathbf{p}_{i \setminus 1}(x)}{\mathbf{p}_{1 \setminus i}(x)}$$

and homomorphically computes $2T + 1$ encrypted evaluations $\left(j, \llbracket \frac{e_{i,j}}{e'_{i,1}} \rrbracket\right)$ for $j \in [2T + 1]$. Using these encrypted evaluations, P_1 homomorphically computes an encrypted rational polynomial $\llbracket \tilde{\mathbf{p}}_i^*(x) \rrbracket$ using rational polynomial interpolation. P_1 then homomorphically reconstructs the roots of $\mathbf{p}_{i \setminus 1}(x)$ and $\mathbf{p}_{1 \setminus i}(x)$ from $\tilde{\mathbf{p}}_i^*$ to obtain $\llbracket D_{i,1}^* \rrbracket, \llbracket D_{1,i}^* \rrbracket$. Note that $\tilde{\mathbf{p}}_i^*(x) = \tilde{\mathbf{p}}_i(x)$ if $\tilde{\mathbf{p}}_i(x)$ has degree at most $2T$, in which case $D_{i,1}^* = D_{i,1}$ and $D_{1,i}^* = D_{1,i}$.

In the final protocol, P_1 homomorphically computes encrypted predicates b_i where $b_i = 1$ iff $\tilde{\mathbf{p}}_i^*(z) = \frac{e'_i}{e'_{i,1}}$ for each $i \in [n] \setminus \{1\}$ and encrypted predicate b' where $b' = 1$ iff $\left|\bigcup_{i \in [n] \setminus \{1\}} (D_{1,i}^* \cup D_{i,1}^*)\right| \leq T$. The output predicate b is homomorphically computed as $\llbracket b \rrbracket = \llbracket b' \cdot \prod_{i \in [n] \setminus \{1\}} b_i \rrbracket$ and jointly decrypted by all the parties. The protocol is formally described in Fig. 6.

Parameters: Parties P_1, \dots, P_n . Each party has a set of m elements. Threshold $T \in \mathbb{N}$. \mathbb{F} is a finite field where $|\mathbb{F}| = \Omega(2^\lambda)$.

Inputs: Party P_i has an input set $S_i = \{a_1^i, \dots, a_m^i\}$ where $a_j^i \in \mathbb{F}$ for all $j \in [m]$.

Output: Each party P_i receives **similar** if $|\left(\bigcup_{i=1}^n S_i\right) \setminus I| \leq T$ and **different** otherwise where $I = \bigcap_{i=1}^n S_i$.

Protocol:

1. Each party P_i generates $(\mathbf{pk}_i, \mathbf{sk}_i) \leftarrow \text{TFHE.DistSetup}(1^\lambda, i)$ and sends \mathbf{pk}_i to P_1 . Then P_1 sends $\mathbf{pk} = (\mathbf{pk}_1 \parallel \dots \parallel \mathbf{pk}_n)$ to all the other parties.
2. P_1 picks a random value $z \in \mathbb{F}$ and sends it to all parties.
3. Each party P_i does the following:
 - (a) Define the polynomial $\mathbf{p}_i(x) := \prod_{a \in S_i} (x - a)$.
 - (b) Compute $e_{i,j} := \mathbf{p}_i(j)$ for $j \in [2T + 1]$ and $e'_i := \mathbf{p}_i(z)$.
 - (c) Send encrypted evaluations $\llbracket e_{i,j} \rrbracket := \text{TFHE.Enc}(\mathbf{pk}, e_{i,j})$ for $j \in [2T + 1]$ and $\llbracket e'_i \rrbracket := \text{TFHE.Enc}(\mathbf{pk}, e'_i)$ to P_1 .
4. P_1 does the following:
 - (a) For each $i \in [n] \setminus \{1\}$, use the algorithm TFHE.Eval to homomorphically compute an encryption $\llbracket \tilde{\mathbf{p}}_i^*(x) \rrbracket$ by rational polynomial interpolation from $2T + 1$ encrypted evaluations $\left(j, \llbracket \frac{e_{i,j}}{e_{1,j}} \rrbracket\right)$ for $j \in [2T + 1]$.
 - (b) For each $i \in [n] \setminus \{1\}$, homomorphically compute the encrypted predicate $\llbracket b_i \rrbracket$ where $b_i = 1$ if $\tilde{\mathbf{p}}_i^*(z) = \frac{e'_i}{e'_1}$ and 0 otherwise.
 - (c) For each $i \in [n] \setminus \{1\}$, homomorphically compute the encrypted roots $\llbracket D_{i,1}^* \rrbracket, \llbracket D_{1,i}^* \rrbracket$ of the numerator and denominator of $\tilde{\mathbf{p}}_i^*(x)$, respectively.
 - (d) Homomorphically compute the encrypted predicate $\llbracket b' \rrbracket$ where $b' = 1$ if $\left|\bigcup_{i \in [n] \setminus \{1\}} (D_{1,i}^* \cup D_{i,1}^*)\right| \leq T$ and 0 otherwise.
5. P_1 sends $\llbracket b \rrbracket = \llbracket b' \cdot \prod_{i \in [n] \setminus \{1\}} b_i \rrbracket$ to all parties who respond with $\llbracket b : \mathbf{sk}_i \rrbracket := \text{TFHE.PartialDec}(\mathbf{sk}_i, \llbracket b \rrbracket)$. P_1 broadcasts $b := \text{TFHE.Combine}(\mathbf{pk}, \{\llbracket b : \mathbf{sk}_i \rrbracket\}_{i \in [n]})$ and all parties output **similar** if $b = 1$ and **different** otherwise.

Fig. 6. Multi-party private intersection cardinality testing protocol $\Pi_{\text{TFHE-C-Test-diff}}$ for $\mathcal{F}_{\text{C-Test-diff}}$

Theorem 3. Assuming threshold FHE with distributed setup, protocol $\Pi_{\text{TFHE-C-Test-diff}}$ (Fig. 6) securely realizes $\mathcal{F}_{\text{C-Test-diff}}$ (Fig. 4).

Proof. Correctness. We first prove the protocol is correct. By the correctness of the TFHE scheme, we only need to show that the computed predicate $b = 1$ if and only if $|\left(\bigcup_{i=1}^n S_i\right) \setminus I| \leq T$. First consider the case where the protocol should output similar. Since

$$\tilde{\mathbf{p}}_i(x) = \frac{\mathbf{p}_i(x)}{\mathbf{p}_1(x)} = \frac{\mathbf{p}_{i \setminus 1}(x)}{\mathbf{p}_{1 \setminus i}(x)},$$

both the numerator and denominator have degree at most T and therefore the rational polynomial interpolation requires at most $(2T + 1)$ evaluation points. Hence $\tilde{\mathbf{p}}_i^*(x) = \tilde{\mathbf{p}}_i(x)$ and $\tilde{\mathbf{p}}_i^*(z) = \tilde{\mathbf{p}}_i(z) = \frac{e'_i}{e'_1}$, thus $b_i = 1$. Since the roots of $\mathbf{p}_{i \setminus 1}$ is simply the set difference $D_{i,1} = S_i \setminus S_1$, we have $D_{i,1}^* = D_{i,1} = S_i \setminus S_1$.

Similarly $D_{1,i}^* = S_1 \setminus S_i$. Since $\left| \bigcup_{i \in [n] \setminus \{1\}} (D_{1,i}^* \cup D_{i,1}^*) \right| = |(\bigcup_{i=1}^n S_i) \setminus I| \leq T$, we have $b' = 1$. Hence the protocol will output $b = 1$.

Now consider the case where the protocol should output different, namely $|(\bigcup_{i=1}^n S_i) \setminus I| > T$. There are two possible cases. In the first case, $|S_i \setminus S_1| > T$ for some i . Then \tilde{p}_i has degree at least $2T + 2$ but \tilde{p}_i^* is interpolated from $2T + 1$ evaluation points, hence $b'_i = 0$ with all but negligible probability. In the second case, $|S_i \setminus S_1| \leq T$ for all $i \in [n] \setminus \{1\}$. Then $D_{i,1}^* = D_{i,1} = S_i \setminus S_1$, $D_{1,i}^* = S_1 \setminus S_i$, and $b_i = 1$ for all i . Since $|(\bigcup_{i=1}^n S_i) \setminus I| > T$, $b' = 0$. In both cases, we have $b = b' \cdot \prod_{i \in [n] \setminus \{1\}} b_i = 0$ with all but negligible probability.

Communication Cost. Each party sends $(2T + 2)$ TFHE encryptions and one partial decryption to P_1 where each plaintext is a field element. P_1 sends one ciphertext to every other party. The size of each encryption and each partial decryption is $\text{poly}(\lambda)$. Thus, the overall communication complexity is $O(n \cdot T \cdot \text{poly}(\lambda))$ in a star network and the protocols runs in $O(1)$ rounds.

Security. The proof of security is identical to the proof of Theorem 2. We defer the formal proof to the full version.

6 TAHE-Based Protocol for $\mathcal{F}_{\text{CTest-diff}}$

In this section, we present a multi-party protocol for private intersection cardinality testing for functionality $\mathcal{F}_{\text{CTest-diff}}$ based on threshold additive homomorphic encryption with distributed setup. That is, the parties learn whether their sets satisfy $|(\bigcup_{i=1}^n S_i) \setminus I| \leq T$. Our protocol works in the star network communication model where P_1 is the central party.

In our construction, we need a secure multi-party computation (MPC) protocol that tests the singularity of a specific Hankel matrix (defined later), which we discuss in Sect. 6.1. Using this, we present our complete protocol in Sect. 6.2.

6.1 Singularity Testing of Hankel Matrices

In Sect. 6.2, we will see that intersection cardinality testing can be reduced to determining whether the determinant of a specific matrix is 0 or not. The latter problem can be reduced to computing the so-called ‘‘Half-GCD’’ of two specific polynomials. In this section, we present a summary of the various results that go into these reductions and refer the reader to the cited works for further details.

Half-GCD Problem. Consider the ring of polynomials $\mathbb{F}[x]$. Note that since $\mathbb{F}[x]$ is a Euclidean domain, Euclid’s GCD algorithm can be applied to polynomials as well. Consider $\mathbf{p}_0, \mathbf{p}_1 \in \mathbb{F}[x]$ with $d = \deg(\mathbf{p}_0) > \deg(\mathbf{p}_1) \geq 0$. The Euclidean algorithm can be viewed as a sequence of transformations of 2-vectors as below:

$$\begin{pmatrix} \mathbf{p}_0 \\ \mathbf{p}_1 \end{pmatrix} \xrightarrow{M_1} \begin{pmatrix} \mathbf{p}_1 \\ \mathbf{p}_2 \end{pmatrix} \xrightarrow{M_2} \dots \xrightarrow{M_{h-1}} \begin{pmatrix} \mathbf{p}_{h-1} \\ \mathbf{p}_h \end{pmatrix} \xrightarrow{M_h} \begin{pmatrix} \mathbf{p}_h \\ 0 \end{pmatrix} \tag{2}$$

Here, M_1, \dots, M_h are 2×2 matrices, $\mathbf{p}_2, \dots, \mathbf{p}_h \in \mathbb{F}[x]$. For vectors U, V and a matrix M , we write $U \xrightarrow{M} V$ to denote $U = MV$.

Equation 2 can be correctly interpreted if we define

$$M_i = \begin{pmatrix} \mathbf{q}_i & 1 \\ 1 & 0 \end{pmatrix}.$$

We call such matrices *elementary matrices*, where \mathbf{q}_i is a polynomial of positive degree. We also refer to \mathbf{q}_i as the *partial quotient* in M_i . A *regular matrix* M is a product of zero or more elementary matrices, namely

$$M = M_1 M_2 \dots M_k \quad (k \geq 0)$$

where if $k = 0$, then M is defined to be the identity matrix of order 2.

We define the *half-GCD (HGCD)* problem for the polynomial ring $\mathbb{F}[x]$ as follows. Given $\mathbf{p}_0, \mathbf{p}_1 \in \mathbb{F}[x]$ with $d = \deg(\mathbf{p}_0) > \deg(\mathbf{p}_1) \geq 0$, compute a regular matrix

$$M = \text{HGCD}(\mathbf{p}_0, \mathbf{p}_1)$$

such that if

$$\begin{pmatrix} \mathbf{p}_0 \\ \mathbf{p}_1 \end{pmatrix} \xrightarrow{M} \begin{pmatrix} \mathbf{p}_2 \\ \mathbf{p}_3 \end{pmatrix},$$

then

$$\deg(\mathbf{p}_2) \geq d/2 > \deg(\mathbf{p}_3).$$

We now recall the result of Thull and Yap [TY90] on the computational complexity of HGCD.

Imported Theorem 4. *Consider the polynomial ring $\mathbb{F}[x]$ and the polynomials $\mathbf{p}_0, \mathbf{p}_1 \in \mathbb{F}[x]$ with $d = \deg(\mathbf{p}_0) > \deg(\mathbf{p}_1) \geq 0$. The computational complexity of the HGCD problem is $O(d \log^2 d)$.*

Singularity Testing of Hankel Matrices. Next, we proceed to outline the results that enable us to use the HGCD problem to test singularity of Hankel matrices. A Hankel matrix is a matrix in which each ascending skew-diagonal from left to right is constant. We will be working with square Hankel matrices. In particular, a $(k + 1) \times (k + 1)$ Hankel matrix takes the form

$$H = \begin{pmatrix} a_0 & a_1 & \dots & a_k \\ a_1 & a_2 & \dots & a_{k+1} \\ \vdots & \vdots & \vdots & \vdots \\ a_k & a_{k+1} & \dots & a_{2k} \end{pmatrix}$$

where the $2k + 1$ entries a_0, a_1, \dots, a_{2k} define H . Define the two polynomials

$$\begin{aligned} \mathbf{p}_0(x) &= x^{2k+1} \\ \mathbf{p}_1(x) &= a_0 + a_1x + a_2x^2 + \dots + a_{2k}x^{2k} \end{aligned}$$

where $\mathbf{p}_0, \mathbf{p}_1 \in \mathbb{F}[x]$. Let $M = \text{HGCD}(\mathbf{p}_0, \mathbf{p}_1)$ and

$$\begin{pmatrix} \mathbf{p}_0 \\ \mathbf{p}_1 \end{pmatrix} \xrightarrow{M} \begin{pmatrix} \mathbf{p}_2 \\ \mathbf{p}_3 \end{pmatrix}.$$

Then we have

$$\deg(\mathbf{p}_2) \geq k + 1 > \deg(\mathbf{p}_3).$$

We recall the setting and results of Brent, Gustavson and Yun [BGY80] that elegantly connect the singularity of H with the HGCD of $\mathbf{p}_0(x)$ and $\mathbf{p}_1(x)$.

Imported Theorem 5. *The Hankel matrix H is singular iff $\deg(\mathbf{p}_3) < k$.*

Putting Imported Theorems 4 and 5 together, we have the following theorem.

Imported Theorem 6. *The computational complexity of testing singularity of a $(k + 1) \times (k + 1)$ Hankel matrix is $O(k \log^2 k)$.*

Multi-party Singularity Testing. Looking ahead, in our multi-party intersection cardinality testing protocol, we will need to test for the singularity of a Hankel matrix H which the parties have additive shares of, and the parties will run a secure multi-party computation (MPC) protocol to jointly test for the singularity of H . The ideal functionality $\mathcal{F}_{\text{SingTest}}$ for the multi-party minimal polynomial computation is defined in Fig. 7. We will need an MPC protocol that realizes $\mathcal{F}_{\text{SingTest}}$ with communication complexity at most $\tilde{O}(k \cdot n \cdot \text{poly}(\lambda))$. Any such protocol suffices, and we denote by Π_{SingTest} the MPC protocol realizing $\mathcal{F}_{\text{SingTest}}$.

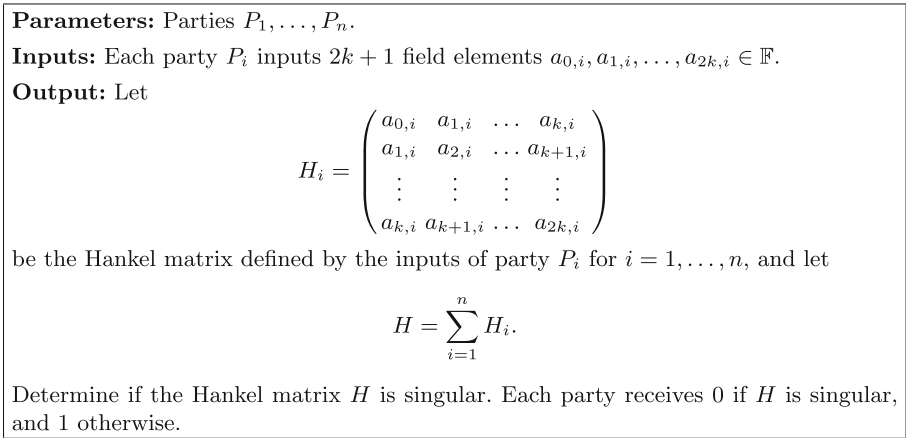


Fig. 7. Ideal functionality $\mathcal{F}_{\text{SingTest}}$ for multi-party singularity testing of a Hankel matrix.

Here we describe two such protocols with communication complexity $\tilde{O}(k \cdot n \cdot \text{poly}(\lambda))$ based on TAHE. In the first protocol, after the TAHE setup, each party

P_i sends $\llbracket H_i \rrbracket$ to P_1 and P_1 homomorphically computes $\llbracket H \rrbracket$. Afterwards P_1 can homomorphically evaluate a circuit C that computes a predicate $b \stackrel{?}{=} (\det(H) = 0)$, following the ideas from [FH96, CDN01]. Finally the parties jointly decrypt the encrypted output. Since the size and depth of C are both $O(k \log^2 k)$ by Imported Theorem 6, the total communication complexity of this protocol is $O(k \log^2 k \cdot n \cdot \text{poly}(\lambda))$ and the round complexity is $O(k \log^2 k)$.

As a second protocol, the parties jointly compute another C' that takes H and a random PRF key r as input and outputs a Yao's garbled circuit [Yao86] that computes C . This approach is inspired by the work of Damgård et al. [DIK+08]. Since both H and r are additively shared among all the parties, this MPC can be done similarly as in the previous protocol, namely P_1 first obtains $\llbracket H \rrbracket$ and $\llbracket r \rrbracket$ and then homomorphically evaluates C' . Since the size C' is $\tilde{O}(k \cdot \text{poly}(\lambda))$ and the depth of C' is constant assuming PRG is a circuit in NC^1 [AIK05], the total communication complexity of this protocol is $\tilde{O}(k \cdot n \cdot \text{poly}(\lambda))$ and the round complexity is $O(1)$.

Two-Party Case. Notice that for two parties, $\mathcal{F}_{\text{SingTest}}$ can be instantiated via Yao's garbled circuits with communication complexity $\tilde{O}(k \cdot \text{poly}(\lambda))$.

6.2 Our Protocol

In this section we present our multi-party private intersection cardinality testing protocol. That is, the parties learn whether their sets satisfy $|(\bigcup_{i=1}^n S_i) \setminus I| \leq T$.

At a high level, our protocol first encodes each party P_i 's set as a polynomial $\mathbf{p}_i(x) = \sum_{j=1}^m x^{a_j^i}$, and let $\mathbf{p}(x) := (n-1)\mathbf{p}_1(x) - \sum_{i=2}^n \mathbf{p}_i(x)$. Notice that a term x^a is cancelled out in the polynomial \mathbf{p} if and only if the element a is in the set intersection I . Therefore, the number of monomials in \mathbf{p} is exactly $|(\bigcup_{i=1}^n S_i) \setminus I|$.

To determine if the number of monomials in \mathbf{p} is $\leq T$, we can apply the polynomial sparsity test of Grigorescu et al. [GJR10] similarly as in [GS19a]. In particular, pick a field \mathbb{F}_q , sample $u \stackrel{\$}{\leftarrow} \mathbb{F}_q$ uniformly at random, and compute the Hankel matrix

$$H = \begin{bmatrix} \mathbf{p}(u^0) & \mathbf{p}(u^1) & \dots & \mathbf{p}(u^T) \\ \mathbf{p}(u^1) & \mathbf{p}(u^2) & \dots & \mathbf{p}(u^{T+1}) \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{p}(u^T) & \mathbf{p}(u^{T+1}) & \dots & \mathbf{p}(u^{2T}) \end{bmatrix}.$$

Determining if the number of monomials in \mathbf{p} is $\leq T$ can be reduced to testing the singularity of H . In particular, we take the following theorem from [GJR10, Theorem 3] and [GS19a, Theorem 1].

Imported Theorem 7. *Let $q > T(T+1)(p-1)2^\kappa$ be a prime. If the number of monomials in \mathbf{p} is $\leq T$, then $\Pr[\det(H) = 0] = 1$, and if the number of monomials in \mathbf{p} is $> T$, then $\Pr[\det(H) = 0] \leq 2^{-\kappa}$,*

1. **Computing Shares of Hankel Matrix H .**
 - (a) P_1 picks a uniform random $u \xleftarrow{\$} \mathbb{F}_q$ and sends to all other parties.
 - (b) P_1 sets a polynomial $\mathbf{p}_1(x) = \sum_{j=1}^m (n-1) \cdot x^{a_j^1}$ in $\mathbb{F}_q[x]$.
 - (c) Each party P_i ($i = 2, 3, \dots, n$) sets a polynomial $\mathbf{p}_i(x) = -\sum_{j=1}^m x^{a_j^i}$ in $\mathbb{F}_q[x]$.
 - (d) Each party P_i ($i = 1, 2, 3, \dots, n$) computes the values $a_{j,i} = \mathbf{p}_i(u^j)$ for $j = 0, 1, \dots, 2T$.
2. **Matrix Singularity Testing of H .** Parties invoke an instance of $\mathcal{F}_{\text{SingTest}}$ where each party P_i inputs $a_{0,i}, \dots, a_{2T,i}$ and obtains a bit b .
3. **Output.** Each party P_i outputs similar if $b = 0$ and different otherwise.

Fig. 8. Multi-party private intersection cardinality testing protocol $\Pi_{\text{CTest-diff}}$.

In our multi-party private intersection cardinality testing protocol, the parties will first compute additive shares of H and then run a multi-party minimal polynomial computation protocol to jointly test the singularity of H . The protocol is presented in Fig. 8.

Theorem 8. *Let $q > T(T+1)(p-1)2^{\kappa}$ be a prime. Assuming threshold additive homomorphic encryption scheme with distributed setup, the protocol $\Pi_{\text{CTest-diff}}$ (Fig. 8) securely realizes $\mathcal{F}_{\text{CTest-diff}}$ in the $\mathcal{F}_{\text{SingTest}}$ -hybrid model.*

Proof. Correctness. By the correctness of $\mathcal{F}_{\text{SingTest}}$, in Step 2 all the parties learn a bit b and $b = 0$ if and only if H is singular, where H is the Hankel matrix $H = \sum_{i=1}^n H_i$ and each Hankel matrix H_i is defined by the inputs of party P_i as

$$H_i = \begin{pmatrix} a_{0,i} & a_{1,i} & \dots & a_{T,i} \\ a_{1,i} & a_{2,i} & \dots & a_{T+1,i} \\ \vdots & \vdots & \ddots & \vdots \\ a_{T,i} & a_{T+1,i} & \dots & a_{2T,i} \end{pmatrix}$$

for $i = 1, \dots, n$. By Imported Theorem 7, $b = 0$ if and only if $|\left(\bigcup_{i=1}^n S_i\right) \setminus I| \leq T$ with all but negligible probability. Therefore the protocol is correct with all but negligible probability.

Communication Cost. The communication cost is the same as the protocol Π_{SingTest} . In particular, the round complexity is $O(1)$ in a star network and the total communication complexity is $\tilde{O}(T \cdot n \cdot \text{poly}(\lambda))$.

Security. We construct a PPT Sim which simulates the view of the corrupted parties. The simulator Sim gets the output $w \in \{\text{similar}, \text{different}\}$ from the ideal functionality. Sim sets a bit $b^* = 1$ if $w = \text{similar}$ and $b^* = 0$ otherwise. Also, for each corrupt party P_i , Sim has as input the tuple (S_i, r_i) indicating the party's input and randomness for the protocol. The strategy of the simulator Sim for our protocol is described below.

1. Invoke the corrupted parties with their corresponding inputs and randomness.

2. Play the role of the honest parties as follows: Run the protocol honestly. Note that P_1 is the only party that ever sends a message, so this step in the simulation is trivial.
3. In Step 2, play the role of $\mathcal{F}_{\text{SingTest}}$ and respond b^* .
4. Finally, output the view of the corrupted parties.

Next we argue that the view of the corrupted parties generated by Sim is computationally indistinguishable to their view in the real world from \mathcal{Z} 's point of view. The only difference between the real and ideal worlds is that in the ideal world, the output from $\mathcal{F}_{\text{SingTest}}$ is replaced by 0 if $|(\bigcup_{i=1}^n S_i) \setminus I| \leq T$ and 1 otherwise. This is computationally indistinguishable from the real world because of the correctness of the protocol.

Corollary 2. *Assuming TAHE with distributed setup, protocol $\Pi_{\text{CTest-diff}}$ (Fig. 8) securely realizes $\mathcal{F}_{\text{CTest-diff}}$ in the star network communication model with communication complexity $\tilde{O}(n \cdot T \cdot \text{poly}(\lambda))$ and round complexity $O(1)$.*

7 Threshold PSI for $\mathcal{F}_{\text{TPSI-diff}}$

Recall that in a multi-party threshold PSI protocol for functionality $\mathcal{F}_{\text{TPSI-diff}}$ defined in Fig. 2, each party wishes to learn the intersection of all their sets if $|(\bigcup_{i=1}^n S_i) \setminus I| \leq T$, that is, if the size of the union of all their sets minus the intersection is less than the threshold T . In this section, we describe our multi-party threshold PSI protocol based on any protocol for multi-party private intersection cardinality testing. We rely on TAHE with distributed setup.

Theorem 9. *Assuming threshold additive homomorphic encryption with distributed setup, protocol $\Pi_{\text{TPSI-diff}}$ (Fig. 9) securely realizes $\mathcal{F}_{\text{TPSI-diff}}$ in the $\mathcal{F}_{\text{CTest-diff}}$ -hybrid model in the star network communication model. Our protocol is secure against a semi-honest adversary that can corrupt up to $(n - 1)$ parties.*

The protocol runs in a constant number of rounds and the communication complexity is $O(n \cdot T \cdot \text{poly}(\lambda))$ in the $\mathcal{F}_{\text{CTest-diff}}$ -hybrid model. We then instantiate the $\mathcal{F}_{\text{CTest-diff}}$ -hybrid with the two protocols from the previous sections: one based on TFHE from Sect. 5.2 that has round complexity $O(1)$ and $O(n \cdot T \cdot \text{poly}(\lambda))$ communication complexity and the other based on TAHE from Sect. 6 that has round complexity $O(1)$ and communication complexity $\tilde{O}(n \cdot T \cdot \text{poly}(\lambda))$. Formally, we get the following corollaries:

Corollary 3. *Assuming TFHE (resp. TAHE) with distributed setup, protocol $\Pi_{\text{TPSI-diff}}$ (Fig. 9) securely realizes $\mathcal{F}_{\text{TPSI-diff}}$ in the star network communication model with communication complexity $O(n \cdot T \cdot \text{poly}(\lambda))$ (resp. $\tilde{O}(n \cdot T \cdot \text{poly}(\lambda))$) and round complexity $O(1)$.*

Our threshold PSI protocol for functionality $\mathcal{F}_{\text{TPSI-int}}$ is almost identical and we defer the details to the full version.

7.1 Protocol

Consider n parties P_1, \dots, P_n with input sets S_1, \dots, S_n of size m and a star network where the central party is P_1 . The parties first run the private intersection cardinality testing protocols for functionality $\mathcal{F}_{\text{CTest-diff}}$ from the previous sections and proceed if $|(\bigcup_{i=1}^n S_i) \setminus I| \leq T$. Then, each party P_i encodes its set as a polynomial $\mathbf{p}'_i(x) = (x - r_i) \cdot \prod_{j=1}^m (x - a_j^i)$ where r_i is picked uniformly at random. The parties then compute $(3T + 4)$ evaluations of the following polynomial $\mathbf{V}(\cdot)$ on points $1, \dots, (3T + 4)$ using threshold additive homomorphic encryption: $\mathbf{V}(x) = \sum_{i=1}^n (\mathbf{p}'_i(x) \cdot \mathbf{R}_i(x))$ where each $\mathbf{R}_i(\cdot)$ is a uniformly random polynomial of degree T that is computed as an addition of n random polynomials - one generated by each party. Then, each party P_i interpolates the degree $(3T + 3)$ rational polynomial $\frac{\mathbf{V}(\cdot)}{\mathbf{p}'_i(\cdot)}$ using the $(3T + 4)$ evaluations. Finally, each party outputs the intersection as $S_i \setminus D_i$ where D_i denotes the roots of the above interpolated polynomial. Our protocol is formally described in Fig. 9.

Two-Party Case. For two parties Alice and Bob, we can rely on AHE alone, where Alice holds the secret key. In particular, define $\mathbf{V}(x) := \mathbf{p}_A(x) \cdot (\mathbf{R}_1^A(x) + \mathbf{R}_1^B(x)) + \mathbf{p}_B(x) \cdot (\mathbf{R}_2^A(x) + \mathbf{R}_2^B(x))$, where $(\mathbf{R}_1^A, \mathbf{R}_2^A)$ and $(\mathbf{R}_1^B, \mathbf{R}_2^B)$ are uniformly random polynomials of degree T generated by Alice and Bob, respectively. To obtain an evaluation of $\mathbf{V}(x)$, Alice first sends an encryption of $\mathbf{p}_A(x)$ and $\mathbf{R}_2^A(x)$ to Bob. Then Bob homomorphically computes an encryption of $r = \mathbf{p}_A(x) \cdot \mathbf{R}_1^B(x) + \mathbf{p}_B(x) \cdot (\mathbf{R}_2^A(x) + \mathbf{R}_2^B(x))$ and sends it back. Alice can decrypt $\llbracket r \rrbracket$ and compute $\mathbf{V}(x) = \mathbf{p}_A(x) \cdot \mathbf{R}_1^A(x) + r$. The communication complexity is $O(T \cdot \text{poly}(\lambda))$.

7.2 Security Proof

Correctness. If $|(\bigcup_{i=1}^n S_i) \setminus I| > T$, then the protocol terminates after the first step - private intersection cardinality testing. If, on the other hand, $|(\bigcup_{i=1}^n S_i) \setminus I| \leq T$, observe that polynomial $\mathbf{V}(x)$ can be rewritten as $\sum_{i=1}^n \mathbf{p}'_i(x) \cdot U_i(x)$ where each U_i is a uniformly random polynomial of degree at most $T + 1$. Now, from the correctness of the TAHE scheme, each party P_i learns $3T + 4$ evaluations of the rational polynomial:

$$\mathbf{q}_i(x) = \frac{\mathbf{V}(x)}{\mathbf{p}'_i(x)} = \frac{\sum_{i=1}^n \mathbf{p}'_i(x) \cdot U_i(x)}{\mathbf{p}'_i(x)} = \frac{\sum_{i=1}^n \mathbf{p}_{i \setminus I}(x) \cdot (x - r_i) \cdot U_i(x)}{\mathbf{p}_{i \setminus I}(x) \cdot (x - r_i)}.$$

Since $|S_i - I| \leq T$ for each $i \in [n]$, the numerator is a polynomial of degree at most $2T + 2$ and the denominator is a polynomial of degree at most $T + 1$. Further, since each U_i is uniformly random, we can show that the numerator is a random degree $2T + 2$ polynomial, and that the gcd of the polynomials in the numerator and denominator is 1 and hence no other terms will get canceled out. The algebraic proofs are deferred to the full version. Therefore, each party P_i can interpolate this rational polynomial using $3T + 4$ evaluation points and thereby learn the numerator and denominator. Finally, observe that for each party P_i ,

Parameters: Parties P_1, \dots, P_n . Each party has a set of m elements. Threshold $T \in \mathbb{N}$. \mathbb{F} is a finite field where $|\mathbb{F}| = \Omega(2^\lambda)$.

Inputs: Party P_i has an input set $S_i = \{a_1^i, \dots, a_m^i\}$ where $a_j^i \in \mathbb{F}$ for all $j \in [m]$.

Output: Each party P_i receives the set intersection $I = \bigcap_{i=1}^n S_i$ if and only if $|\left(\bigcup_{i=1}^n S_i\right) \setminus I| \leq T$.

Protocol:

1. **Private Intersection Cardinality Testing.** The parties invoke $\mathcal{F}_{\text{CTest-diff}}$ on inputs S_1, \dots, S_n and receive back $w \in \{\text{similar}, \text{different}\}$. If $w = \text{different}$ then all parties output \perp .
2. **TAHE Key Generation.** Each party P_i generates $(\text{pk}_i, \text{sk}_i) \leftarrow \text{TAHE.DistSetup}(1^\lambda, i)$ and sends pk_i to P_1 . Then P_1 sends $\text{pk} = (\text{pk}_1 \| \dots \| \text{pk}_n)$ to all the other parties.
3. **Evaluations of Random Polynomial.** In this phase the parties will evaluate a polynomial

$$V(x) = \sum_{i=1}^n \left(\mathbf{p}'_i(x) \cdot \left(R_1(x) + \dots + R_{i-1}(x) + \tilde{R}_i(x) + R_{i+1}(x) \dots + R_n(x) \right) \right)$$
 for $x \in [3T + 4]$ where the terms are defined as follows.
 - (a) Each party P_i defines $\mathbf{p}_i(x) = \prod_{j=1}^m (x - a_j^i)$ and $\mathbf{p}'_i(x) = \mathbf{p}_i(x) \cdot (x - r_i)$ where $r_i \xleftarrow{\$} \mathbb{F}$.
 - (b) Each party P_i uniformly samples $R_i, \tilde{R}_i \xleftarrow{\$} \mathbb{F}[x]$ of degree $T + 1$, computes $R_i(x)$ for $x \in [3T + 4]$ and sends encrypted $\llbracket R_i(x) \rrbracket$ to P_1 .
 - (c) For each $i \in [n], x \in [3T + 4]$, party P_1 sends $\llbracket e_{i,x} \rrbracket = \left\llbracket \sum_{j \in [n] \setminus \{i\}} R_j(x) \right\rrbracket$ to P_i .
 - (d) For each $x \in [3T + 4]$, each party P_i sends $\llbracket v_{i,x} \rrbracket = \left\llbracket \mathbf{p}'_i(x) \cdot \left(e_{i,x} + \tilde{R}_i(x) \right) \right\rrbracket$ to P_1 .
 - (e) For each $x \in [3T + 4]$, P_1 sends $\llbracket v_x \rrbracket = \left\llbracket \sum_{i=1}^n v_{i,x} \right\rrbracket$ to all P_i .
 - (f) For each $x \in [3T + 4]$, each party P_i sends $\llbracket v_x : \text{sk}_i \rrbracket \leftarrow \text{TAHE.PartialDec}(\text{sk}_i, \llbracket v_x \rrbracket)$ to P_1 .
 - (g) For each $x \in [3T + 4]$, P_1 sends $V(x) = \text{TAHE.Combine}(\text{pk}, \{\llbracket v_x : \text{sk}_i \rrbracket\}_{i \in [n]})$ to all P_i .
4. **Computing Set Intersection.** Each party P_i does the following:
 - (a) Interpolate $\mathbf{q}_i(x)$ to be the degree $3T + 3$ rational polynomial such that $\mathbf{q}_i(x) = \frac{V(x)}{\mathbf{p}'_i(x)}$ for $x \in [3T + 4]$ and the gcd of the numerator and denominator is 1. Let D_i be the roots of the denominator of $\mathbf{q}_i(x)$.
 - (b) Output the set intersection $I = S_i \setminus D_i$.

Fig. 9. Multi-party threshold PSI protocol $\Pi_{\text{TPSI-diff}}$ for functionality $\mathcal{F}_{\text{TPSI-diff}}$.

the roots of the denominator contains the set $S_i \setminus I$ and a random r_i , from which P_i can easily compute the intersection I .

Communication Cost. The first phase of the protocol, namely private intersection cardinality testing, has a communication complexity of $O(n \cdot T \cdot \text{poly}(\lambda))$ when instantiated with the TFEH-based scheme in Sect. 5.2 and a communication complexity of $\tilde{O}(n \cdot T \cdot \text{poly}(\lambda))$ when instantiated with the TAHE-based scheme in Sect. 6.

We now analyze the communication cost for the second phase where the parties compute the concrete intersection. The TAHE key generation is independent of the set sizes and the threshold T and has a communication complexity of only $O(n \cdot \text{poly}(\lambda))$. The bottleneck of the protocol is in Step 3, that is, evaluating the random polynomial. In Steps 3b, 3d, and 3f, every party sends $3T+4$ encryptions or partial decryptions to P_1 hence the cost for these steps is $O(n \cdot T \cdot \text{poly}(\lambda))$. In Steps 3c, 3e, and 3g, P_1 sends $3T+4$ ciphertexts or plaintexts to every other party so the cost of these steps is $O(n \cdot T \cdot \text{poly}(\lambda))$. Finally, the last stage, namely computing the set intersection, does not involve any communication. Thus, the overall communication cost for computing the intersection is $O(n \cdot T \cdot \text{poly}(\lambda))$.

Therefore, when the private intersection cardinality testing protocol is instantiated with the TFHE-based protocol, the overall communication complexity is $O(n \cdot T \cdot \text{poly}(\lambda))$ and when instantiated with the TAHE-based scheme, the overall communication complexity is $\tilde{O}(n \cdot T \cdot \text{poly}(\lambda))$ for some a priori fixed polynomial $\text{poly}(\cdot)$ and is independent of the size of each input set m .

Security. Consider an environment \mathcal{Z} who corrupts a set \mathcal{S}^* of n^* parties where $n^* < n$. The simulator Sim has the output of the functionality $\mathcal{F}_{\text{TPSI-diff}}$, namely the intersection set I or \perp . Sim sets $w = \text{similar}$ if the output is I and $w = \text{different}$ if the output is \perp . In addition, Sim has the tuple (S_i, r_i) for each corrupt party P_i indicating the party's input and randomness for the protocol. The strategy of the simulator Sim for our multi-party threshold PSI protocol is described below.

(a) Private Intersection Cardinality Testing: Sim plays the role of the ideal functionality $\mathcal{F}_{\text{CTest-diff}}$ and responds with w .

(b) TAHE Key Generation: Sim runs the distributed key generation algorithm $\text{TAHE.DistSetup}(1^\lambda, i)$ of the TAHE scheme honestly on behalf of each honest party P_i as in the real world. Note that Sim also knows $(\{\text{sk}_i\}_{i \in \mathcal{S}^*})$ as it knows the randomness for the corrupt parties.

(c) Evaluations of Random Polynomial: Sim does the following:

1. Encode the intersection set $I = \{b_1, \dots, b_{|I|}\}$ as a polynomial as follows: $\text{p}_I(x) = \prod_{i=1}^{|I|} (x - b_i)$.
2. Pick a random polynomial $U(\cdot)$ of degree $2T+2$ and set the polynomial $\text{V}(x)$ as follows: $\text{V}(x) = \text{p}_I(x) \cdot U(x)$.
3. In Steps 3b–3e, on behalf of every honest party P_i , whenever P_i has to send any ciphertext, send $\llbracket 0 \rrbracket$ using fresh randomness.
4. For each $x \in [3T+4]$, let $\llbracket v_x \rrbracket$ denote the ciphertext that is sent to all the parties at the end of Step 3f.
5. In Step 3f, for each $j \in [3T+4]$, on behalf of each honest party P_i , instead of computing $\{\llbracket v_x : \text{sk}_i \rrbracket\}$ by running the honest TAHE.PartialDec algorithm as in the real world, Sim computes the partial decryptions by running the simulator TAHE.Sim as follows: $\{\llbracket v_x : \text{sk}_i \rrbracket\} \leftarrow \text{TAHE.Sim}(\text{C}, \text{V}(x), \llbracket v_x \rrbracket, \{\text{sk}_i\}_{i \in \mathcal{S}^*})$, where C is the public linear circuit to compute $\text{V}(x)$ by P_1 .
6. Finally, in Step 3g, if P_1 is honest, send the evaluations of polynomial $\text{V}(x)$ as in the real world description.

Hybrids. We now show that the above simulation strategy is successful against all environments \mathcal{Z} that corrupt parties in a semi-honest manner. That is, the view of the corrupt parties along with the output of the honest parties is computationally indistinguishable in the real and ideal worlds. We will show this via a series of computationally indistinguishable hybrids where the first hybrid Hybrid_0 corresponds to the real world and the last hybrid Hybrid_4 corresponds to the ideal world.

- **Hybrid₀ - Real World:** In this hybrid, consider a simulator SimHyb that plays the role of the honest parties as in the real world.
- **Hybrid₁ - Private Intersection Cardinality Testing:** In this hybrid, SimHyb plays the role of the ideal functionality $\mathcal{F}_{\text{CTest-diff}}$ and responds with similar if $|\left(\bigcup_{i=1}^n S_i\right) \setminus I| \leq T$ and different otherwise.
- **Hybrid₂ - Simulate Partial Decryptions:** In this hybrid, in the evaluations of random polynomial, SimHyb simulates the partial decryptions generated by the honest parties in Step 3f as done in the ideal world. That is, for each $j \in [3T + 4]$, SimHyb computes the partial decryptions as $\{\llbracket v_x : \text{sk}_i \rrbracket\} \leftarrow \text{TAHE.Sim}(\mathbb{C}, \mathbb{V}(x), \llbracket v_x \rrbracket, \{\text{sk}_i\}_{i \in \mathcal{S}^*})$. Observe that the polynomial $\mathbb{V}(\cdot)$ is still computed as in the real world (and in Hybrid_2).
- **Hybrid₃ - Switch Polynomial Computation:** In this hybrid, the polynomial $\mathbb{V}(\cdot)$ is no longer computed as in the real world. Instead, SimHyb now picks a random polynomial $U(\cdot)$ of degree $2T + 2$ and sets the polynomial $\mathbb{V}(\cdot)$ as follows: $\mathbb{V}(x) = \mathbf{p}_I(x) \cdot U(x)$.
- **Hybrid₄ - Switch Encryptions:** In this hybrid, in the evaluations of random polynomial, SimHyb now computes every ciphertext generated on behalf of any honest party as encryptions of 0 as done by Sim in the ideal world. This hybrid corresponds to the ideal world.

We show that every pair of consecutive hybrids is computationally indistinguishable in the full version.

References

- [AIK05] Applebaum, B., Ishai, Y., Kushilevitz, E.: Computationally private randomizing polynomials and their applications. In: CCC (2005)
- [BDP21] Branco, P., Döttling, N., Pu, S.: Multiparty cardinality testing for threshold private set intersection. In: PKC (2021)
- [Ben94] Benaloh, J.: Dense probabilistic encryption May 1994
- [BFK+19] Badrinarayanan, S., Fernando, R., Koppula, V., Sahai, A., Waters, B.: Output compression, MPC, and IO for turing machines. In: ASIACRYPT (2019)
- [BGG+18] Boneh, D., Gennaro, R., Goldfeder, S., Jain, A., Kim, S., Rasmussen, P.M.R., Sahai, A.: Threshold cryptosystems from threshold fully homomorphic encryption. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018. LNCS, vol. 10991, pp. 565–596. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-96884-1_19

- [BGY80] Brent, R.P., Gustavson, F.G., Yun, D.Y.Y.: Fast solution of Toeplitz systems of equations and computation of padé approximants. *J. Algorithms* **1**(3), 259–295 (1980)
- [BO15] Braverman, M., Oshman, R.: On information complexity in the broadcast model. In: *PODC* (2015)
- [BPSW07] Brickell, J., Porter, D.E., Shmatikov, V., Witchel, E.: Privacy-preserving remote diagnostics. In: *CCS* (2007)
- [CDN01] Cramer, R., Damgård, I., Nielsen, J.B.: Multiparty computation from threshold homomorphic encryption. In: Pfitzmann, B. (ed.) *EUROCRYPT* 2001. LNCS, vol. 2045, pp. 280–300. Springer, Heidelberg (2001). https://doi.org/10.1007/3-540-44987-6_18
- [CM20] Chase, M., Miao, P.: Private set intersection in the internet setting from lightweight oblivious PRF. In: Micciancio, D., Ristenpart, T. (eds.) *CRYPTO* 2020. LNCS, vol. 12172, pp. 34–63. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-56877-1_2
- [DCT10] De Cristofaro, E., Tsudik, G.: Practical private set intersection protocols with linear complexity. In: *FC* (2010)
- [DCW13] Dong, C., Chen, L., Wen, Z.: When private set intersection meets big data: an efficient and scalable protocol. In: *CCS* (2013)
- [DIK+08] Damgård, I., Ishai, Y., Krøigaard, M., Nielsen, J.B., Smith, A.: Scalable multiparty computation with nearly optimal work and resilience. In: Wagner, D. (ed.) *CRYPTO* 2008. LNCS, vol. 5157, pp. 241–261. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-85174-5_14
- [FH96] Franklin, M., Haber, S.: Joint encryption and message-efficient secure computation. *J. Cryptol.* **9**, 217–232 (1996)
- [FNP04] Freedman, M.J., Nissim, K., Pinkas, B.: Efficient private matching and set intersection. In: Cachin, C., Camenisch, J.L. (eds.) *EUROCRYPT* 2004. LNCS, vol. 3027, pp. 1–19. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24676-3_1
- [GJR10] Grigorescu, E., Jung, K., Rubinfeld, R.: A local decision test for sparse polynomials. *Inf. Process. Lett.* **110**(20), 898–901 (2010)
- [GN19] Ghosh, S., Nilges, T.: An algebraic approach to maliciously secure private set intersection. In: Ishai, Y., Rijmen, V. (eds.) *EUROCRYPT* 2019. LNCS, vol. 11478, pp. 154–185. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-17659-4_6
- [GS19a] Ghosh, S., Simkin, M.: The communication complexity of threshold private set intersection. In: Boldyreva, A., Micciancio, D. (eds.) *CRYPTO* 2019. LNCS, vol. 11693, pp. 3–29. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26951-7_1
- [GS19b] Ghosh, S., Simkin, M.: The communication complexity of threshold private set intersection (2019). ia.cr/2019/175
- [HFH99] Huberman, B.A., Franklin, M., Hogg, T.: Enhancing privacy and trust in electronic communities. In: *Proceedings of the 1st ACM Conference on Electronic Commerce* (1999)
- [HOS17] Per, A., Hallgren, C.O., Sabelfeld, A.: Privacy-preserving ridesharing. In: *CSF, Privatepool* (2017)
- [HV17] Hazay, C., Venkatasubramanian, M.: Scalable multi-party private set-intersection. In: *PKC* (2017)
- [HW15] Hubáček, P., Wichs, D.: On the communication complexity of secure function evaluation with long output. In: *ITCS* (2015)

- [IKN+17] Ion, M., et al.: Private intersection-sum protocol with applications to attributing aggregate ad conversions (2017). ia.cr/2017/738
- [KKRT16] Kolesnikov, V., Kumaresan, R., Rosulek, M., Trieu, N.: Efficient batched oblivious PRF with applications to private set intersection. In: CCS (2016)
- [KMP+17] Kolesnikov, V., Matania, N., Pinkas, B., Rosulek, M., Trieu, N.: Practical multi-party private set intersection from symmetric-key techniques. In: CCS (2017)
- [KMWF07] Kiltz, E., Mohassel, P., Weinreb, E., Franklin, M.: Secure linear algebra using linearly recurrent sequences. In: TCC (2007)
- [KS05] Kissner, L., Song, D.: Privacy-preserving set operations. In: Shoup, V. (ed.) CRYPTO 2005. LNCS, vol. 3621, pp. 241–257. Springer, Heidelberg (2005). https://doi.org/10.1007/11535218_15
- [MPR+20] Miao, P., Patel, S., Raykova, M., Seth, K., Yung, M.: Two-sided malicious security for private intersection-sum with cardinality. In: Micciancio, D., Ristenpart, T. (eds.) CRYPTO 2020. LNCS, vol. 12172, pp. 3–33. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-56877-1_1
- [MTZ03] Minsky, Y., Trachtenberg, A., Zippel, R.: Set reconciliation with nearly optimal communication complexity. IEEE Trans. Inf. Theory **49**, 2213–2218 (2003)
- [MZ17] Mohassel, P., Zhang, Y.: SecureML: a system for scalable privacy-preserving machine learning. In: IEEE S and P (2017)
- [NMH+10] Nagaraja, S., Mittal, P., Hong, C.-Y., Caesar, M., Borisov, N.: BotGrep: finding P2P bots with structured graph analysis. In: USENIX security symposium (2010)
- [OOS16] Orrù, M., Orsini, E., Scholl, P.: Actively secure 1-out-of-n OT extension with application to private set intersection. In: CT-RSA (2016)
- [Pai99] Paillier, P.: Public-key cryptosystems based on composite degree residuosity classes. In: Stern, J. (ed.) EUROCRYPT 1999. LNCS, vol. 1592, pp. 223–238. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-48910-X_16
- [PRTY19] Pinkas, B., Rosulek, M., Trieu, N., Yanai, A.: SpOT-light: lightweight private set intersection from sparse OT extension. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019. LNCS, vol. 11694, pp. 401–431. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-26954-8_13
- [PRTY20] Pinkas, B., Rosulek, M., Trieu, N., Yanai, A.: PSI from PaXoS: fast, malicious private set intersection. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020. LNCS, vol. 12106, pp. 739–767. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-45724-2_25
- [PSSZ15] Pinkas, B., Schneider, T., Segev, G., Zohner, M.: Phasing: private set intersection using permutation-based hashing. In: USENIX (2015)
- [PSWW18] Pinkas, B., Schneider, T., Weinert, C., Wieder, U.: Efficient circuit-based PSI via cuckoo hashing. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018. LNCS, vol. 10822, pp. 125–157. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-78372-7_5
- [PSZ14] Pinkas, B., Schneider, T., Zohner, M.: Faster private set intersection based on OT extension. In: USENIX (2014)
- [RR17] Rindal, P., Rosulek, M.: Malicious-secure private set intersection via dual execution. In: CCS (2017)
- [TPKC07] Troncoso-Pastoriza, J.R., Katzenbeisser, S., Celik, M.: Privacy preserving error resilient dna searching through oblivious automata. In: CCS (2007)

- [TY90] Thull, K., Yap, C.: A unified approach to HGCD algorithms for polynomials and integers. Manuscript (1990)
- [Yao86] Yao, A.C.-C.: How to generate and exchange secrets (extended abstract). In: FOCS (1986)
- [ZC18] Zhao, Y., Chow, S.S.M.: Can you find the one for me? Privacy-preserving matchmaking via threshold PSI (2018). ia.cr/2018/184